Benchmark Dose Software (BMDS)

VERSION 3.0
USER GUIDE

⊕EPA

Authors and Reviewers

Federal Authors

Jeff Gift, Ph.D.
J Allen Davis, M.S.
Todd Blessinger, Ph.D.
U.S. EPA
Office of Research and Development
National Center for Environmental Assessment (NCEA)
RTP, NC, Cincinnati, OH and Washington, DC

Matthew Wheeler
The National Institute for Occupational Safety and Health (NIOSH)

Contract Authors

Louis Olszyk
Cody Simmons
Bruce Allen
Michael E. Brown
General Dynamics Information Technology
US EPA, 109 T.W. Alexander Dr.
Research Triangle Park, NC 27711

Reviewers

Laura Carlson
EPA National Center for Environmental Assessment

Stephen Gilbert
The National Institute for Occupational Safety and Health (NIOSH)

Contents

Table of Tables

Table of Figures

EPA

Abbreviations

AIC	Akaike's information criterion
BIC	Bayesian Information Criterion
BMD	benchmark dose
BMDL	benchmark dose lower confidence limit
BMDS	Benchmark Dose Software
BMDU	benchmark dose upper confidence limit
BMR	benchmark response
CI	confidence interval
CSF	cancer slope factor
CV	coefficient of variation
EPA	Environmental Protection Agency
IRIS	Integrated Risk Information System
LOAEL	lowest-observed-adverse-effect level
LPP	Log Posterior Probability
LSC	Litter specific covariate
MLE	maximum-likelihood estimation
NCEA	National Center for Environmental Assessment
NCTR	National Center for Toxicological Research
NIEHS	National Institute of Environmental Health Sciences
NIOSH	National Institute for Occupational Safety and Health
NOAEL	no-observed-adverse-effect level
POD	point of departure
RfC	inhalation reference concentration
RfD	oral reference dose
SD	standard deviation
SE	standard error
SEM	standard error of the mean
US EPA	United States Environmental Protection Agency

1.0   Overview

The U.S. Environmental Protection Agency (EPA) Benchmark Dose Software (BMDS) was developed as a tool to facilitate the application of benchmark dose (BMD) methods to EPA hazardous pollutant risk assessments. This user guide provides instruction on how to use the BMDS, but is not intended to address or replace EPA BMD guidance. However, every attempt has been made to make this software consistent with EPA guidance, including the Risk Assessment Forum (RAF) Benchmark Dose Technical Guidance Document (U.S. EPA, 2012).

1.1   History of BMDS Development

Research into model development for BMDS started in 1995 and the first BMDS prototype was internally reviewed by EPA in 1997. After external and public reviews in 1998-1999, and extensive Quality Assurance testing in 1999-2000, the first public version of BMDS, version 1.2, was released in April 2000.

A complete history of the versions of BMDS released by EPA is contained in the Version History appendix. The models contained in the current version of BMDS are listed in Section 9.0. "Model Descriptions." on page 40.

1.2   How EPA Uses BMD Methods

EPA uses BMD methods to derive risk estimates such as reference doses (RfDs), reference concentrations (RfCs), and Cancer Slope Factors (CSF), which are used along with other scientific information to set standards for human health effects.

Prior to the availability of tools such as BMDS, noncancer risk assessment benchmarks such as RfDs and RfCs were determined from no-observed-adverse-effect levels (NOAELs), which represent the highest experimental dose for which no adverse health effects have been documented.

However, using the NOAEL in determining RfDs and RfCs has long been recognized as having limitations:

- It is limited to one of the doses in the study and is dependent on study design
- It does not account for variability in the estimate of the dose-response
- It does not account for the slope of the dose-response curve
- It cannot be applied when there is no NOAEL, except through the application of an uncertainty factor (Kimmel and Gaylor, 1988; Crump, 1984).

A goal of the BMD approach is to define a starting point of departure (POD) for the computation of a reference value (RfD or RfC) or cancer slope factor (CSF) that is more independent of study design. The EPA Risk Assessment Forum has published technical guidance for the application of the BMD approach in cancer and non-cancer dose-response assessments (U.S. EPA, 2012).

Using BMD methods involve fitting mathematical models to dose-response data and using the different results to select a BMD that is associated with a predetermined

benchmark response (BMR), such as a 10% increase in the incidence of a particular lesion or a 10% decrease in body weight gain.

BMDS facilitates these operations by providing simple data-management tools and an easy-to-use interface to run multiple models on the same dose-response dataset. Results from all models include model run options chosen by the user, goodness-of-fit information, the BMD, and the estimate of the lower-bound confidence limit on the BMD (BMDL). Model results are presented as textual and graphical outputs that can be printed or saved and incorporated into other documents.

1.3   Future of BMDS

EPA plans to continually improve and expand the BMDS system. Plans include developing an online version of BMDS that will be integrated with the EPA Health & Environmental Research Online (HERO) database and Health Assessment Workspace Collaborative (HAWC) website, adding covariate analysis tools and Bayesian models and model averaging methods for continuous response data to further alleviate issues and uncertainties associated with data selection, bounding frequentist model parameters, and assisting the user with selecting a "best" model.

Use the BMDS web page as your most up-to-date source of information and updates pertaining to the BMDS. The entire BMDS system or model updates can be downloaded from the web site. The source code files for the models used in the BMDS system are also available via the BMDS web site to reviewers and programmers who might be interested in performing an in-depth analysis of the model algorithms and features.

We welcome and encourage your comments on the BMDS software. Please provide comments, recommendations, suggested revisions, or corrections through our Help Desk Form.

2.0   What's New in BMDS 3.0

BMDS 3.0 is a major re-design of BMDS that contains substantial model code and interface enhancements that reflect nearly two decades of experience and feedback on the needs of risk assessors with respect to benchmark dose modeling.

New Bayesian dichotomous models and a Bayesian dichotomous model averaging feature have been added. Pre-existing dichotomous and continuous models have been upgraded with new features and recoded to stabilize and improve performance.

2.1   Interface Enhancements

2.1.1 Analysis Workbook

Datasets, modeling and reporting options for an analysis are entered in the BMDS 3.0 Analysis Workbook. Users can specify modeling options from intuitive forms and picklists. All calculations are performed within the Analysis Workbook.

The Analysis Workbook is designed to facilitate performing and tracking dose-response analyses of multiple dichotomous, combined tumor, nested dichotomous, and continuous response datasets. Depending on the needs of the risk assessment, users can focus a BMDS 3.0 analysis on datasets associated by:

- study (e.g., for chemicals with a large database of studies)
- chemical (e.g., for chemicals that are not well-studied)
- health outcome (e.g., for chemicals with health outcomes that have been assessed in multiple studies and/or by multiple response measures)

2.1.2 Results Workbook

All datasets, modeling and reporting options entered in the BMDS 3.0 Analysis Workbook can be saved and retrieved at any time prior to or after modeling.

When modeling is performed, all model results are recorded in a separate Results Workbook for each dataset analyzed.

All the options used in the analysis are saved in each Results Workbook such that the analysis can be re-initiated in the BMDS 3.0 Analysis Workbook from the Results Workbook (e.g., to rerun the analysis using different options).

2.2   New Model Additions

New to the BMDS model suite are Bayesian versions of all traditional frequentist dichotomous models and a Bayesian model averaging feature (currently only available for dichotomous data).

Details of the new Bayesian models and modeling averaging feature are provided in Section 9.6, "Bayesian Dichotomous Model Descriptions," on page 54.

2.3 Upgrades to Pre-Existing Models

BMDS continuous models have been upgraded to include a "Hybrid" modeling capability and Lognormal response options. For more details, refer to Section 5.2, "Continuous Model Options," on page 19 and Section 9.4, "Frequentist Continuous Model Descriptions," on page 42.

Also, all pre-existing models have been re-coded to facilitate their maintenance and improve their performance in terms of stability, accuracy, reliability, and speed.

Note   BMDS 3.0 handles Akaike Information Criterion (AIC) calculations somewhat differently from BMDS 2.x to facilitate comparing models with different likelihoods (i.e., normal vs. lognormal). For more details, refer to Section 7.3, "AIC for Continuous Models," on page 26.

2.4 Backwards-Compatibility

BMDS 3.0 retains backwards compatibility with BMDS 2.7 and BMDS Wizard 1.11. For experienced users, BMDS 3.0 resembles the pre-existing BMDS Wizard in these ways:

- Excel-based
- Enables users to see and specify modeling options in a single worksheet
- Includes auto-selection features for identifying the "best" results in accordance with EPA recommendations or user-defined logic
- Documents all inputs and outputs in a single results workbook for each dataset modeled
- Provides flexible print options for displaying results in Microsoft Word tables formatted in a manner suitable for presentation in a risk assessment

2.5 Models Not Included in BMDS 3.0

BMDS 3.0 contains all the models and features that were available in BMDS 2.7 and BMDS Wizard 1.11 except for:

- Dichotomous background dose models
- Rai and Van Ryzin nested dichotomous models
- ToxicoDiffusion model
- Ten Berge model

These models are not being maintained or supported by EPA at this time.

However, these models can be accessed in BMDS 2.7, which is available from the BMDS website as an archive version of BMDS.

The Ten Berge model is superseded by the latest version of EPA's categorical regression software (CatReg), which has the same functionality but with added features and options.

3.0   Setting Up BMDS 3.0

3.1   System Requirements

- BMDS 3.0 is distributed as a .zip file on the BMDS website download page, which can be unzipped to any folder where the user has read/write privileges (Administrator privileges are not required).
- BMDS 3.0 requires Microsoft Excel 2016 or later with macros enabled (visit the Microsoft support site for information on enabling Excel macros).

Note   Check the BMDS website periodically for updates to the software or help manual.

3.2   Creating a BMDS Desktop Icon

You may find it more convenient to run BMDS from a desktop shortcut icon. To do so:

1. Delete any older BMDS shortcut icons on your desktop.
2. In Windows Explorer, navigate to the newly installed BMDS application folder.
3. Right-click the BMDSxxx.xlsm file (where "xxx" denotes the current BMDS version number). A context menu appears.
4. Click Send To. A submenu appears.
5. Click "Desktop (Create Shortcut)". Windows creates a shortcut to the file on your desktop.

3.3   Uninstalling Previous Versions of BMDS

It is not necessary to uninstall previous versions of BMDS to install and run BMDS 3.0.

4.0 Tutorial: Analyzing Multiple Datasets in BMDS 3.0

BMDS 3.0 has been designed to facilitate the process of analyzing, recording, and reporting dose-response analyses in a manner that is typically necessary and consistent with EPA recommendations and guidelines for the development of an EPA chemical risk assessment.

Once a dataset(s) has been entered and modeling options defined the BMDS 3.0 analysis can be saved and recalled at any time.

To perform and save a BMDS analysis a user must:

1. Document the analysis
2. Add datasets
3. Select and save modeling options
4. Run the models, review the results, and prepare a summary report (if desired)

The following sections provide a simplified tutorial that overviews each step of the process. The tutorial also references more detailed explanations located elsewhere in this documentation.

4.1   Step 1: Analysis Documentation

The first step in performing an analysis is to identify a directory where you will store your Results Workbooks (dataset-specific modeling results) and Word Report files (created from dataset-specific modeling results). This is done in the Analysis Workbook of BMDS.

Analysis specifications can be saved as either:

- **Save Analysis Without Running.** If this button is clicked, a Save Workbook with no results will be created, with a name corresponding to the user-specified (or default) "Analysis Name" entered on the "Main" worksheet of the Analysis Workbook, and saved to the user-specified output directory.

- **Run Analysis.** If this button is clicked after all steps of an analysis are completed, dataset-specific Results Workbook and Word Report files with names corresponding to the dataset names provided on the "Data" worksheet (Step 2) will be automatically generated with user-specified modeling options (Step 3) and saved to the user-specified output directory.

All Analysis Workbook specifications are saved in the Save Workbook and Results Workbook files so their corresponding Analysis Workbook can be re-generated. When the **Load Analysis** button is selected from within a Save Workbook and the Results Workbook file, the Analysis Workbook will be re-generated with the saved specifications, including all modeling options and datasets.

Note   When an Analysis Workbook is re-generated, all options and datasets in an open Analysis Workbook will be overwritten.

You can also provide an Analysis Description in the "Main" worksheet of the Analysis Workbook. While not a required step, such documentation is useful for analyses that you want to save for future use or consideration.

As stated in Section 2.1.1. "Analysis Workbook." on page 9, BMDS 3.0 has been designed such that the saved "analysis" can be a collection of dose-response datasets from a single study, all dose-response datasets available for a chemical (e.g., for

chemicals that are not well-studied), or all dose-response datasets related to a particular health outcome (e.g., multiple measures of cardiovascular effects).

4.2 Step 2: Add Datasets

After entering the analysis documentation information in the "Main" tab, the dose-response data should be entered in the "Data" worksheet of BMDS.

The user can add multiple datasets associated with four response types:

- summarized continuous (e.g., mean and SD)
- individual continuous (e.g., dose and response for each test subject)
- dichotomous (e.g., lesion incidence)
- nested dichotomous (e.g., developmental study) responses.

In each case, dose-response data can be entered as integers or integers with decimals (e.g., fractions).

To insert new datasets, the user must choose a response type and identify the number of rows required.

To paste data that was copied from a table or spreadsheet in another program (e.g., a prior version of BMDS), you may need to right-click on the destination cell and select "Paste Special" from the Excel context menu. An import feature that will allow users to import datasets directly from .dax files created by prior BMDS versions will be added in a future version.

Enter a unique name for each dataset. This name will be used by BMDS to reference the dataset on the "Main" worksheet (where users can select datasets to include in a modeling analysis) and to name all Result Workbook and Word Report files generated from modeling the dataset. The user can also add more detailed notes to describe the dataset and these notes can be displayed in the Results Workbook and Word Report file.

4.2.1 For Dichotomous Response Data

The default column headers are "Dose," "N" and "Incidence" (grey row), but for reporting purposes the user can enter replacement terms (blue row) such as "mg/kg-day," "Subjects" and "Cases."

4.2.2 For Continuous Response Data

For summarized continuous response data, the default column headers are "Dose," "N," "Mean" and "Std. Dev." (grey row).

For individual continuous response data, the default column headers are "Dose," "Response."

Again, for reporting purposes the user can enter replacement terms (blue row). The BMDS "Data" worksheet offers a tool for converting standard errors to the required standard deviation metric.

For continuous response data, the user will be given the choice (on the "Main" worksheet) to either allow BMDS to choose the adverse direction based on the dose-response trend, or manually identify the dose-response direction as "Up" or "Down." This will impact the derivation of the benchmark response (BMR) for which the benchmark dose (BMD) is estimated.

4.2.3 For Nested Dichotomous Data

The default column headers are "Dose," "Litter Size," "Incidence" and "Litter Specific Covariate" (LSC) (grey row). Again, for reporting purposes the user can enter their own replacement terms (blue row).

There must be data in the LSC row even if the modeling options do not call for the use of LSC.

4.3 Step 3: Select and Save Modeling Options

All models and modeling options available for use in an analysis can be selected in the "Main" worksheet of BMDS. Options can be saved and reloaded at any time before or after running an analysis. An analysis can involve the use of any or all of four Model Types:

- Continuous
- Dichotomous
- Dichotomous – Multi-tumor
- Dichotomous – Nested

Each model type offers a different set of models and/or modeling options. For more details, refer to Section 5.0, "Modeling Options," on page 19 and Section 9.0, "Model Descriptions," on page 40.

As in previous versions of BMDS, users can choose to run multiple models in an analysis. However, unlike previous versions of BMDS, BMDS 3.0 allows users to run the selected models against multiple, user-defined modeling "Option Sets" and multiple datasets. The BMDS 3.0 "Main" worksheet lists the datasets entered in the "Data" worksheet, enabling the user to choose datasets of the appropriate Model Type to analyze using the selected Models and Option Sets. The results for each Model-Option Set combination are recorded in separate worksheets within dataset-specific Results Workbooks.

4.3.1 Continuous Response Models and Options

All the traditional frequentist models and options that were available for analyzing continuous response data in previous versions of BMDS are available in BMDS 3.0 (Bayesian models and Bayesian model averaging will be available for continuous responses in a future version of BMDS).

Also, users are now able to use the Hybrid continuous modeling method and the lognormal response distribution assumption (previously only available for Exponential models) for all continuous models.

As in previous versions of BMDS, the user can choose to run the Hill, Polynomial, and Power models restricted or unrestricted; the Linear model is not restricted and the Exponential models can only be run restricted.

In the "Main" worksheet of the BMDS 3.0 Analysis Workbook, the user can define multiple Option Sets to apply to multiple user-selected models and multiple user-selected datasets in a single "batch" process. The adverse direction of each dataset can be manually set to "Up" (increasing with dose) or "Down" (decreasing with dose) or "Auto-detect" via trend testing (default).

Continuous model Option Sets are user-defined with respect to:

- BMR Type: Standard Deviation, Relative Deviation, Absolute Deviation, Point, Hybrid – Extra Risk
- BMRF: BMR factor
- Tail Probability: Cut-off for defining adversity; only applicable to Hybrid extra risk model
- Confidence Level: fraction between 0 and 1; 0.95 is recommended by EPA (2012)
- Distribution: Normal or Log-normal response distribution
- Variance: Constant or Non-constant
- Polynomial Restrictions: Beta parameter restriction; Automatic (depends on detected or specific adverse direction), Non-negative or Non-positive
- Background: Not currently specifiable for continuous models

For more details on the continuous models and relevant options, refer to:

- Section 5.2. "Continuous Model Options." on page 19
- Section 9.4, "Frequentist Continuous Model Descriptions," on page 42

4.3.2 Dichotomous Response Models and Options

BMDS 3.0 offers the traditional frequentist dichotomous response models available in previous versions of BMDS plus Bayesian versions of each model, as well as a Bayesian model averaging feature.

Most frequentist models can be run restricted or unrestricted. The EPA default recommendation for initial runs is to restrict the Gamma, Log-Logistic, Multistage and Weibull models and un-restrict the Log-Probit and Dichotomous Hill model; the Logistic, Probit and Quantal Linear models are not restricted.

Dichotomous model Option Sets are user-defined with respect to:

- BMR Type: Extra Risk or Added Risk
- BMR: a fraction between 0 and 1; EPA standard is 0.1
- Confidence Level: fraction between 0 and 1; 0.95 is recommended by EPA (2012)
- Background: Estimated, Zero or User-Specified; usually estimated unless strong evidence for zero or specific value

For more details on the dichotomous models and relevant options, refer to:

- Section 5.1. "Dichotomous Model Options." on page 19
- Section 9.5. "Frequentist Dichotomous Model Descriptions." on page 48
- Section 9.6. "Bayesian Dichotomous Model Descriptions." on page 54

4.3.3 Dichotomous—Multi-tumor Models and Options

As in previous versions of BMDS, BMDS 3.0 allows users to run the EPA's Multi-tumor (MS_Combo) model to determine the BMD, BMDL and BMDU that is associated with a benchmark response (BMR) for the risk of experiencing any combination of the multiple tumor types.

However, unlike previous versions of BMDS, BMDS 3.0 provides users with the option to manually select or allow BMDS to "Auto-select" the degree of Multistage model to apply to a dataset. The auto-selection process follows the most recent EPA technical guidance for selecting the Multistage model degree for the analysis of cancer datasets, which differs from the model selection process described by EPA (2012) for other modeling scenarios.

Dichotomous model Option Sets are user-defined with respect to:

- BMR Type: Extra Risk or Added Risk
- BMR: a fraction between 0 and 1; EPA standard is 0.1
- Confidence Level: fraction between 0 and 1; 0.95 is recommended by EPA (2012)
- Background (dataset-specific): Estimated, Zero or User-Specified; usually estimated unless strong evidence for zero or specific value

For more details on the dichotomous multi-tumor models and relevant options, refer to:

- Section 5.1. "Dichotomous Model Options." on page 19
- Section 6.0. "Multiple Tumor Analysis." on page 23
- Section 9.8. "Multi-tumor (MS_Combo) Model Description." on page 63

4.3.4 Dichotomous—Nested Models and Options

BMDS 3.0 allows users to run the EPA's Nested Logistic nested dichotomous model.

> **Note**   The NCTR (National Center for Toxicological Research) nested dichotomous model is not in the BMDS 3.0 release. It will be included in a future release.

EPA no longer supports the Rai and Van Ryzin model that was available in previous versions of BMDS.

Unlike previous versions of BMDS, BMDS 3.0 does not require the user to specify the model form, but rather automatically runs all forms of the available nested models (i.e., all combinations of "Use Litter Specific Covariate." "Do Not Use Litter Specific Covariate," "Estimate Intralitter Correlations" and "Assume Intralitter Correlations of Zero").

Dichotomous nested model Option Sets are user-defined with respect to:

- BMR Type: Extra Risk or Added Risk
- BMR: a fraction between 0 and 1; EPA standard is 0.1
- Confidence Level: fraction between 0 and 1; 0.95 is recommended by EPA (2012)
- Litter Specific Covariate: Overall or Control Group Mean
- Background (dataset-specific): Estimated, Zero or User-Specified; usually estimated unless strong evidence for zero or specific value

For more details on dichotomous nested models and their options, refer to

- Section 5.3. "Nested Model Options." on page 22
- Section 9.7. "Frequentist Nested Model Descriptions." on page 58

4.4   Step 4: Run Models, Review Results, and Prepare Summary Report(s)

When modeling is performed, the results are recorded in separate Results Workbooks for each dataset analyzed. All of the options used in the analysis are saved in each Results Workbook such that the analysis can be re-initiated in the BMDS 3.0 Analysis Workbook from the Results Workbook (e.g., to rerun the analysis using different options).

In the Report Options worksheet of the BMDS 3.0 Analysis Workbook, users can select modeling inputs and results to report for each model type, "Export Options," and "Word Report Options."

- Export Options selected affect both the Result Workbook and Word Report files that are generated from an analysis.

- Word Report Options are applied for the generation of tabular documentation of modeling results in Microsoft Word.
- Export Options ("User Inputs" and "Analysis Results") are set separately for each of the different model/analysis types.

Because the Word report may take a few minutes to compile, we recommend running the Word Report Options *after* you have verified your analysis results. When you are satisfied with the analysis results, then re-run with the Word Report Option checked.

Users familiar with the previous BMDS Wizard will note that BMDS 3.0 uses a similar approach to analyzing modeling results and making automatic recommendations regarding model selection that are consistent with the 2012 EPA Benchmark Dose Technical Guidance (U.S. EPA, 2012). These criteria can be altered in the "Logic" worksheet of the BMDS 3.0 Analysis Workbook, as presented below in Figure 1. Decision logic can be turned on or off, and specific criteria can be enabled or disabled for different dataset types. Notice that the logic depends on what type of data is being analyzed (continuous, dichotomous, nested).

Figure 1. BMDS 3.0 "Logic" worksheet with EPA default recommendation decision logic.

Model Recommendation/Bin Placement Logic

Test Description	Test On/Off			Test Threshold (where appropriate)	Bin Placement if Test is Failed	Notes to Show
	Continuous	Dichotomous	Nested			
BMD is not calculated	On	On	On		Unusable Bin	BMD not estimated
BMDL is not calculated	On	On	On		Unusable Bin	BMDL not estimated
BMDU is not calculated	Off	Off	Off		No Bin Change (Warning)	BMDU not estimated
AIC is not calculated	On	On	On		Unusable Bin	AIC not estimated
Constant Variance	On			0.05	Questionable Bin	Constant variance test failed (Test 2 p-value < 0.05)
Non-Constant Variance	On			0.05	Questionable Bin	Non-constant variance test failed (Test 3 p-value < 0.05)
Goodness of fit p-test	On	On	On	0.1	Questionable Bin	Goodness of fit p-value < 0.1
Goodness of fit p-test (cancer)	Off	On	Off	0.05	Questionable Bin	Goodness of fit p-value < 0.05
Ratio of BMD/BMDL (serious)	On	On	On	20	Questionable Bin	BMD/BMDL ratio > 20
Ratio of BMD/BMDL (caution)	On	On	On	5	No Bin Change (Warning)	BMD/BMDL ratio > 5
Abs(Residual of interest) too large	On	On	On	2	Questionable Bin	|Residual for Dose Group Near BMD| > 2
BMD higher than highest dose	On	On	On	1	No Bin Change (Warning)	BMD higher than maximum dose
BMDL higher than highest dose	On	On	On	1	No Bin Change (Warning)	BMDL higher than maximum dose
BMD lower than lowest dose (warning)	On	On	On	3	No Bin Change (Warning)	BMD 3x lower than lowest non-zero dose
BMDL lower than lowest dose (warning)	On	On	On	3	No Bin Change (Warning)	BMDL 3x lower than lowest non-zero dose
BMD lower than lowest dose (serious)	On	On	On	10	Questionable Bin	BMD 10x lower than lowest non-zero dose
BMDL lower than lowest dose (serious)	On	On	On	10	Questionable Bin	BMDL 10x lower than lowest non-zero dose
Abs(Residual at control) too large	On	On	On	2	No Bin Change (Warning)	|Residual at control| > 2
Poor control dose std. dev.	On			1.5	No Bin Change (Warning)	Modeled control response std. dev. >|1.5| actual response std. dev
D.O.F. equals 0	On	On	On		Questionable Bin	d.f.=0, saturated model (Goodness of fit test cannot be calculated)

Based on the decision logic entered by the user as described above, BMDS will attempt to select a "recommended" model. A user must ultimately select a model and may choose to disagree with the BMDS auto-determination. BMDS 3.0 automatically generates suggested text for the "BMDS Recommendation" and "BMDS Recommendation Notes" columns of the Results Workbook summary tables and the Word Report File tables. While some reformatting is allowed in the Results Workbook (e.g., row heights, column widths, and the size, design, and position of plots), the text and numeric results cannot be modified. However, the Word Report files can be modified extensively, and the user is encouraged to take advantage of this flexibility to change and/or expand on the table headers and the justification provided for why a model was selected.

BMDS 3.0 places each model into one of three different bins:

- Viable—highest quality model, no serious deficiencies found based on user-defined logic but may contain warnings
- Questionable—serious deficiencies with model based on user-defined decision logic
- Unusable—required outputs such as BMD or BMDL are not estimated

After all models with the same BMR have been placed into one of three different quality bins, a model is recommended from the highest quality bin based on BMDL or AIC criteria defined in the 2012 EPA Benchmark Dose Technical Guidance (U.S. EPA, 2012). The default setting for "sufficiently close" BMDLs is a 3-fold range. Figure 2 reflects the BMDS 3.0 model recommendation logic using the default assumptions shown in Figure 1.

Figure 2. *Flow chart of BMDS 3.0 model recommendation logic using EPA default logic assumptions.*
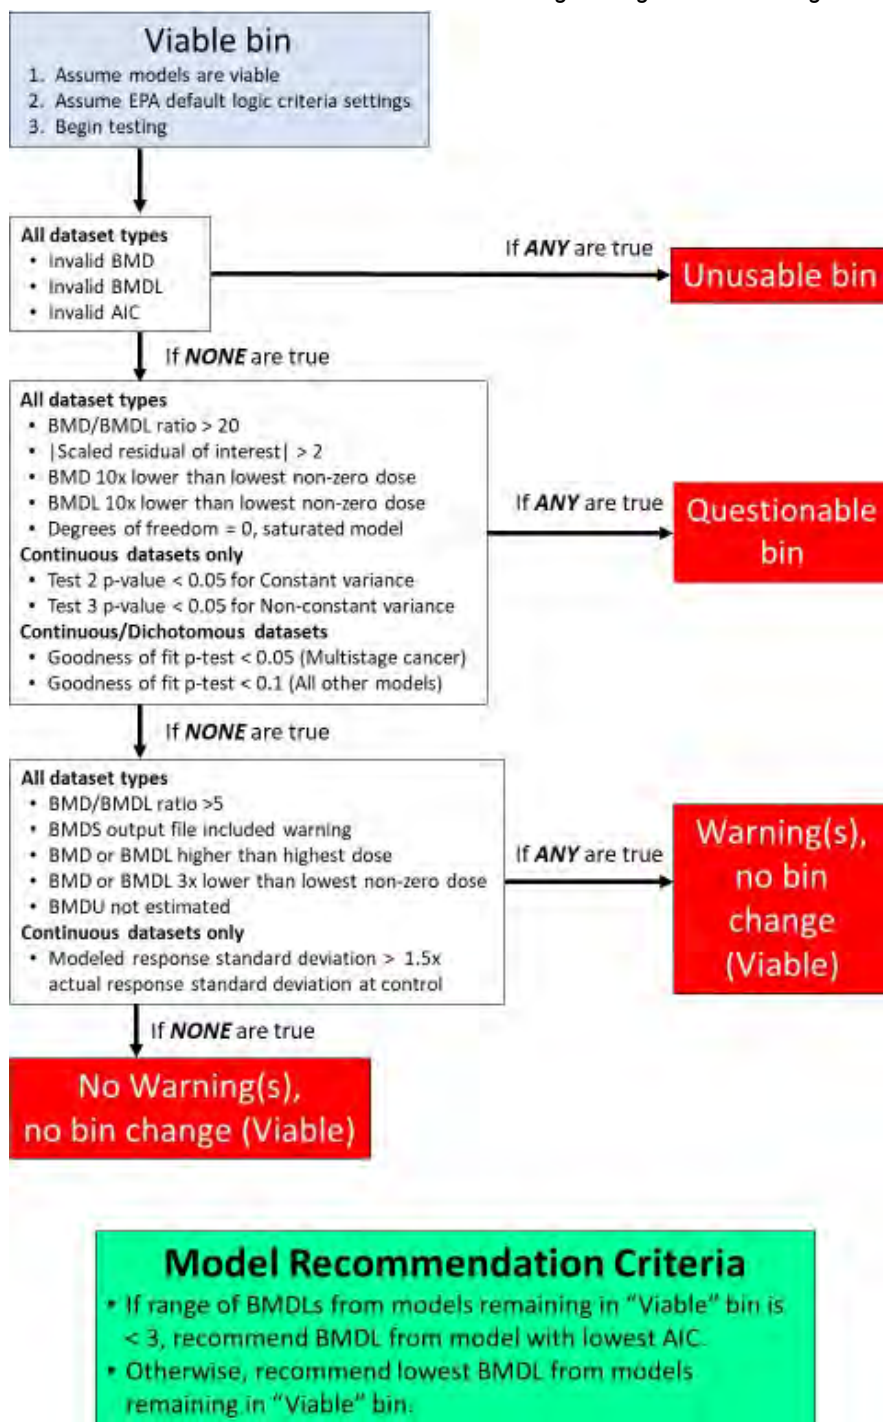
5.0   Modeling Options

5.1   Dichotomous Model Options

Risk Type

Choices are "Extra" (Default) or "Added."

Added risk is the additional proportion of total animals that respond in the presence of the dose, or the predicted probability of response at dose d, $P(d)$, minus the predicted probability of response in the absence of exposure, $P(0)$: $P(d) - P(0)$

Extra risk is the additional risk divided by the predicted proportion of animals that will not respond in the absence of exposure, $1 - P(0)$: $\frac{P(d) - P(0)}{1 - P(0)}$ *The BMRF for all dichotomous models must be between 0 and 1 (not inclusive).*

5.1.1 Note about BMR and Graphs

The response associated with the BMR that is displayed in the graphical model output will only be the same as the BMR when $P(0) = 0$. This is because to obtain the actual response value one must solve for $P(d)$ in the equation for added or extra risk discussed above.

5.2   Continuous Model Options

Constant Variance

When selected (default), the model assumes a constant variance across all dose groups. If not selected, then the model assumes the variance can be different for each dose group, and that the variance varies as a power function of the mean response. For more details, refer to Section 9.4, "Frequentist Continuous Model Descriptions," on page 42.

Adverse Direction

Choices for the Adverse Direction option are "Automatic" (default), "Up," or "Down." This option refers to whether adversity increases as the dose-response curve rises "up" or falls "down." If automatic is chosen, the software chooses the adverse direction based on the shape of the dose-response curve. Manually choose the adverse direction if you know the direction of adversity for the endpoint being studied. This selection only impacts how the user-designated BMR is used in conjunction with model results to obtain the BMD.

BMR Type

The BMR type is the method of choice for defining the response level used to derive the benchmark dose (BMD). The choices allowed are "Rel. Dev." (default), "Abs. Dev.," "Std. Dev.," "Point," and "Hybrid" (Hill model only).

- **Rel. Dev. (Relative Deviation)** means the response associated with the BMR will be the background estimate plus or minus (depending on the Adverse Direction) the product of the background estimate times the BMRF entered by the user.
- **Abs. Dev. (Absolute Deviation)** means the response associated with the BMR will be the background estimate plus or minus the BMRF.

- **Std. Dev. (Standard Deviation)** means the response associated with the BMR will be the background estimate plus or minus the product of the BMRF times the standard deviation for the control group data.
- **Point** means the response associated with the BMR will be the BMRF value itself.
- **Hybrid** Defines the BMD through the adverse probability of response. Here the BMRF represents the increased probability of an adverse response given the BMD.

$Rel.Dev.Response \ = \ \mu(0) \ + \ (BMRF * \mu(0))$ (Default)

$Abs.Dev.Response \ = \ \mu(0) \ + \ BMRF$

$Std.Dev.Response \ = \ \mu(0) \ + \ (BMRF * STD)$

$Point\ Response \ = \ BMRF$

Hybrid

Solution to "up": $BMRF \ = \ \dfrac{\Pr(X > X_0|D) - \Pr(X > X_0|0)}{1 - \Pr(X > X_0|0)}$

Solution to "down": $RBMRF \ = \ \dfrac{\Pr(X < X_0|D) - \Pr(X < X_0|0)}{1 - \Pr(X < X_0|0)}$

where $\Pr(X < X_0 \mid 0)$ is the background probability that defines adverse response.

Note	When response data is lognormally distributed, the BMR Types acquire different meanings. As of BMDS 3.0, all continuous exponential models can assume lognormal distribution.

5.2.1 Lognormal Response Option

When modeling continuous response data, the standard assumption for the BMDS continuous models is that the underlying distributions (one for each dose group) are Normal, with a mean given by the dose-response model and a variance as specified by the user (constant or a function of the mean response). An alternative assumption is that the responses are Lognormally distributed.

In BMDS 3.0 all continuous models allow the user to choose between Normal and Lognormal response distribution assumptions (unlike prior versions of BMDS, which only allowed this for Exponential models). If the user has access to the individual response data, those data can be log-transformed prior to analysis but, as discussed below, this is not a recommended approach. If the user suspects that the responses are Lognormally distributed, the recommended approach is to model the untransformed data assuming the underlying distribution is Lognormal with median values defined by the dose-response function and a constant log-scale variance, corresponding to an assumption of a constant coefficient of variation. For more details, refer to Section 5.2.2, "Definition of BMR Types under Lognormal Distribution Assumption," on page 21.

BMDS 3.0 provides an exact maximum-likelihood estimation (MLE) solution when data are assumed to be lognormally distributed and individual response data are available. When the data are assumed to be lognormally distributed and the data are presented in terms of group-specific means and standard deviations, then the exact MLE solution cannot be obtained. In that case, the "Solution" is "Approximate" and the means and standard deviations of the log-transformed data are estimated as follows:

$$log - scale\ mean \ = \ ln(mean) \ - \ ln(1 + \frac{\left(\frac{std}{mean}\right) \times 2}{2})$$

$$log - scale\ std\ = \sqrt{ln(1 + \left(\frac{std}{mean}\right) \times 2)}$$

Using log-transformed responses in the analysis is not recommended, for the following reasons:

- If you choose to log-transform the data prior to analysis, then the interpretation of the BMD and BMDL estimates would have to be considered carefully (and perhaps in consultation with a statistician). Data interpretation when using log-transformed responses will not be the same as when using the natural-scale response values. Indeed, the models—when "transformed back" to the natural scale—will not correspond to any of the standard BMDS models.

  For example, if using the power model on log-transformed responses, the user is actually implicitly modeling the medians (on the natural-scale) with the function $e^{(background+slope\times dose^{power}}$ which is not a standard BMDS model and whose characteristics (e.g., exponential increases in response) may not be those desired by the user.

- Similarly, the interpretation of the BMD will not correspond to simple expressions (e.g., if the BMR is set equal to a relative deviation of 10%, that relative deviation will be assessed on the log-scale and so will not yield BMD or BMDL estimates that correspond to a 10% change in the original mean responses).

For these reasons, log-transforming the response values is not considered a "best practice" and, as stated, should only be applied and interpreted with supporting statistical expertise. Therefore, in most cases, the user should use non-transformed values and select the lognormal distribution if the data is assumed to be lognormally distributed.

5.2.2 Definition of BMR Types under Lognormal Distribution Assumption

The BMDS continuous models allow the user to assume that the response data are lognormally distributed, with median values defined by the dose-response function and a constant log-scale variance. Under such an assumption the BMR types are defined and implemented so that they are calculated by the program to return BMDs as follows (where BMRF is the numerical value, specified by the user, indicating the response, or change in response, of interest):

- **Relative Deviation**: The natural scale median value at the BMD, $m(BMD)$, differs from the natural scale median at 0 dose, $m(0)$, such that $\frac{|m(BMD) - m(0)|}{m(0)} = BMRF$.

- **Absolute Deviation**: The natural scale median value at the BMD, $m(BMD)$, differs from the natural scale median at 0 dose, $m(0)$, such that $|m(BMD) - m(0)| = BMRF$.

- **Standard deviation**: The log-scale mean at the BMD, $ln(m(BMD))$, differs from the log-scale mean at 0 dose, $ln(m(0))$, such that $\frac{|ln(m(BMD)) - ln(m(0))|}{\sigma(0)} = BMRF$, where $\sigma(0)$ is the log-scale standard deviation at 0 dose. Recall that $\sigma(0) = ln(GSD(0))$. This definition allows the user to use BMRF's typical of an analysis where a normal distribution of responses is assumed (e.g., the EPA default of 1 standard deviation) and still maintain the logic and rationale for such choices, since the log-transformed response values under the lognormal assumption would themselves be normally distributed.

- **Point**: The natural scale median value at the BMD, $m(BMD)$, equals the BMRF, i.e., $m(BMD) = BMRF$.

5.3   Nested Model Options

Risk Type

Choices are "Extra" or "Added." Additional risk is the additional proportion of total animals that respond in the presence of the dose, or the probability of response at dose d, $P(d)$, minus the probability of response in the absence of exposure, $P(0)$: $P(d) - P(0)$. Extra risk is the additional risk divided by the proportion of animals that will not respond in the absence of exposure, $1 - P(0)$: $\frac{P(d) - P(0)}{1 - P(0)}$. Thus, extra and additional risk are equal when background rate is zero.

The following were options in previous versions of BMDS. BMDS 3.0 does not require the user to specify these options anymore, but simply runs and provides output for all possible option combinations.

Use a Litter Specific Covariate

Optionally, enables the user to account for inter-litter variability by using a litter specific covariate (LSC). If the box is checked (default), Theta values are estimated. If the box is unchecked, the Theta values are set to zero. Do not use LSC if the corresponding metric is affected by dose or if its use does not sufficiently improve model fit, as indicated by a lower AIC value.

Fixed LSC Value

Choices are "Control Group Mean" (default) or "Overall Mean." See Section 9.7, "Frequentist Nested Model Descriptions" on page 58 for an explanation as to why this option is necessary, and which choice would be preferred for your given dataset. Basically, the Overall Mean should be used under most circumstances. If the Litter Specific Covariate differs from dose to dose (without any apparent consistent trend with respect to dose), consider using the Control Group Mean.

Intralitter Correlations

Provides user with the option to allow the models to attempt to estimate intralitter correlations or assume they are zero. If "Estimate Intralitter Correlations" is selected (default), all the Phi values are estimated (one for each dose group). If "Assume Intralitter Correlations Zero" is chosen, all the Phi values are set to zero.

6.0   Multiple Tumor Analysis

6.1   Assumptions and Results

The analyses of multiple tumors have the following assumptions and results.

1. The tumors are statistically independent of one another. ***Note:*** Unless there is substantial biological evidence to indicate that the tumor types are ***not*** independent—conditional on model parameter values—the approach based on independence is considered appropriate.

2. A multistage model is an appropriate model for each of the tumors separately. The individual multistage models fit to the individual tumors need not have the same polynomial degree, however.

3. The user is interested in estimating the risk of getting one or more of the tumors being analyzed; the results indicate the BMD and BMDL associated with the user-defined benchmark response (BMR) level, where the BMD and BMDL are the maximum likelihood and lower bound estimates of the dose that is estimated to give an extra risk equal to the BMR for the "combination" (getting one or more of the tumors).

In accordance with EPA cancer guidelines, a Multiple Tumor Analysis will always run the restricted form of the Multistage model. A new feature in BMDS 3.0 allows users to have BMDS "Auto-Select" the appropriate polynomial degree of the Multistage model for each tumor dataset. When the "Auto-Select" feature is used, BMDS runs all relevant forms of the Multistage model and selects the polynomial degree to use based on the current EPA Multistage model selection criteria for tumor analyses. This is the default option in BMDS 3.0, but the user can also choose to manually set the polynomial degree for each dataset. In any case, it is ultimately the user's responsibility to ensure that the degree of the polynomial and other selections for modeling parameters are as desired and appropriate for the dataset(s) being analyzed.

6.2   Obtaining the Combined BMD

Per EPA cancer guidelines, the Multi-tumor model uses the restricted form of the Multistage model. Because of the form of the restricted multistage model, the combined BMD is obtained in a relatively straightforward manner from the maximum likelihood parameter estimates from the models fit to the individual tumors.

The combined maximum log-likelihood is the sum of the individual maximized log-likelihoods (summed over the individual tumor analyses).

The combined BMD is the dose that is estimated to yield an extra risk of getting one or more of the tumors, where the extra risk is equal to the BMR.

The calculation of the combined BMDL is a more complicated computation based on the profile-likelihood approach.

As such, it gives the lowest value of the dose that satisfies the following conditions:

*   There is a combination of parameters (across all models) for which the value of the BMDL gives a combined extra risk equal to the BMR and, using those parameter values,

- The combined log-likelihood is greater than or equal to a minimum log-likelihood defined by the maximum log-likelihood and the confidence level specified by the user (i.e., the parameters that give the desired extra risk when the dose is equal to the BMDL give a combined log-likelihood that is "close enough" to the maximum combined log-likelihood).

The fitted model log-likelihoods for continuous endpoints are reported in the Likelihoods of Interest (continuous endpoints) tables.

6.3 Running an Analysis and Viewing Results

The user should first enter the dichotomous tumor datasets to be analyzed in the Data worksheet of BMDS. Then the dichotomous datasets entered in the Data worksheet will be available and can be selected for use in a Multi-tumor analysis in the Main worksheet of BMDS.

Once the datasets to be analyzed are selected, the user needs to set dataset-specific modeling options for Multistage Degree ("Auto-Select" or specify) and Background ("Estimated", "Zero" or "User-Specified") and general modeling options for Risk Type ("Extra Risk" or "Added Risk"; EPA recommends use of "Extra Risk"), BMR and Confidence Level.

When "Run Analysis" is selected a separate Results Workbook of multi-tumor results is created. The workbook will include results for each individual tumor considered separately (using the chosen dataset-specific options), and the corresponding estimate of the BMD and BMDL for the combined tumor probability for the risk type, BMR and confidence levels specified by the user.

Plots for individual multistage model runs will be shown on the individual model results tabs. If the "Auto-Select" feature was used to select the Multistage polynomial degree, the user should verify that the resultant model fits are adequate in the desired dose-response region. If the user wants to try a different Multistage polynomial degree they can re-run the analysis using a specified degree instead of "Auto-Select."

6.4 Troubleshooting a Tumor Analysis

If one or more of the tumors is estimated to have a BMD greater than three times the highest dose tested (for that tumor), then the multiple tumor analysis will stop at an intermediate point, i.e., after the fitting has been done for the tumor in question and the magnitude of that BMD has been determined. No tumors listed below that tumor will be analyzed and no combination will be completed.

It is probably the case that the tumor in question will not add substantially to the estimation of a BMD for the combinations of tumors, assuming other tumors have BMDs less than three times the highest dose; that is because the magnitude of response for the tumor in question has not even reached the benchmark response level for such a high exposure and so its individual contribution to the risk of getting one or more of the tumors being analyzed will be small in comparison to that for the other tumors. The user might attempt a combination that does not include the tumor in question.

7.0   Other BMD Analysis Considerations

7.1   Continuous Response Data with Negative Means

Data with negative means should only be modeled with a constant variance model.

It may occasionally be the case that, when modeling transformed data, you will need to model negative data. In this case, the transformation used should be a variance-stabilizing transformation so that a constant-variance model would be appropriate.

If a standard deviation-based BMR is used to define the BMD calculations, then a constant can be added to all the observations (or means) to make the values (means) positive. That will not change the standard deviations of the observations and would allow you to model the variance.

7.2   Test for Combining Two Datasets for the Same Endpoint

At this time, BMDS does not include a formal test for similarity of dose response across covariate values (e.g., across class variables like species or sex). EPA's categorical regression software, CatReg, has that capability.

However, the following procedure can be used in BMDS if you have dose-response data for two experiments that you are considering combining (e.g., for the two sexes within a species, or two species, etc.).

1. Choose a single model to consider for both dataset.
2. Model the two runs separately. For each run, record the following:
    - Maximum log-likelihood for each dataset. Add the numbers from each dataset to get the **summed** log-likelihood.
    - The number of unconstrained parameters for each dataset. Add the numbers from each run to get the **summed** unconstrained parameters.
3. Combine the data from the two experiments and model them together. Record the following:
    - The maximum log-likelihood for the combined dataset. This will be the **combined** log-likelihood. The fitted model log-likelihoods are reported in the Analysis of Deviance (dichotomous endpoints) or Likelihoods of Interest (continuous endpoints) tables.
    - The number of unconstrained parameters for the combined dataset. This will be the **combined** unconstrained parameters.
4. Subtract the combined log-likelihood from the summed log-likelihood. Then, multiply the difference by 2.
5. Compare the value from Step 4 to a chi-squared distribution. The degrees of freedom for that chi-squared distribution will be the difference between the summed unconstrained parameters (Step 2) and the combined unconstrained parameters (Step 3).

    If the value from Step 4 is in the tail (say, greater than the 95th percentile) of the chi-squared distribution in question, then reject the null hypothesis that the two sets have the same dose-response relationship. If rejection occurs, then infer that it is not proper to combine the two datasets.

7.3   AIC for Continuous Models

To facilitate comparing models with different likelihoods (i.e., normal vs. lognormal), the log-likelihood for the normal and lognormal distributions are calculated using all normalizing constants. *This results in different numerical AIC values than those given in earlier BMDS versions.*

"Even though the BMDS 3.0 AIC values differ from those in BMDS 2.x versions, if the models have the same underlying distribution, then the difference of the AICs will be the same as previous versions of BMDS. This assumes that the BMDS 3.0 and BMDS 2.x model fits are the same for the two models being compared. The AIC difference may not be the same if one or more of the model fits differ between the two versions (e.g., if one or more of the 3.0 models provide an improved fit to the data over the corresponding BMDS 2.x model).

However, when comparing models having different parametric distributions, the AIC differences will not be the same as previous BMDS versions. For these comparisons, the AIC calculated using the BMDS 3.0 software is correct and will result in the proper comparison between any two models regardless of underlying distribution.

A note of caution is required for situations where only the sufficient statistics are approximated using the lognormal distribution. In these cases, comparisons between models using the normal distribution should not be made using the AIC.

8.0 Output from Individual Models

Dataset-specific Results Workbooks generated by the BMDS 3.0 Analysis Workbook contain separate worksheets for each Model-Option Set combination that consist of tabular and graphical summaries of the modeling inputs and results.

The purpose of these worksheets is to provide the user with goodness-of-fit criteria and model results to aid in determining the appropriateness of the Model and Option Set to the benchmark dose derivation.

This section describes BMDS model outputs that are common to all model types. Succeeding sections describe model-specific outputs.

8.1 Model Run Documentation (User Input Table)

The Results Workbook worksheets that are generated for each Model-Options set contain a User Input table that gives the user a quick verification of the options they had selected for that Model-Options set.

For instance, when two users may be comparing results and obtained different answers, they may consult their respective User Input tables to make sure the settings were the same or if they had used the same (or most current) version of the models.

The User Input tables within the Results Workbook worksheets for each Model-Option Set contain the model name and version number, the dataset name, dataset information, including user notes, entered by the user on the "Data" worksheet of BMDS, and modeling options entered by the user on the "Main" worksheet of BMDS.

8.2 Benchmark Dose Estimates and Key Fit Statistics (Benchmark Dose Table)

The Benchmark Dose table of the Results Workbook worksheets contains the BMD, BMDL, and BMDU estimates, AIC, and the overall goodness of fit test p-value, chi-square, and degrees of freedom (df) for each Model Option set analyzed.

For more information on how the BMD, BMDL and BMDU values are derived, refer to Sections 9.4.3 and 9.4.4.

The Akaike's Information Criterion (AIC) (Akaike, 1973) value given on the BMDS Results Workbook worksheets is -2L + 2p, where L is the log-likelihood at the maximum likelihood estimates for the parameters, and p is the number of model parameters estimated (and not on a restriction boundary). It can be used to compare different types of models which use a similar fitting method (for example, least squares or a binomial maximum likelihood), as do all dichotomous, continuous and nested model types within BMDS. The model with the lowest AIC would be presumed to be the better model under this method. Although such methods are not exact, they can provide useful guidance in model selection.

> ***Note*** BMDS 3.0 handles Akaike Information Criterion (AIC) calculations somewhat differently from BMDS 2.x To facilitate comparing models with different likelihoods (i.e., normal vs. lognormal). For more details, refer to Section 7.3, "AIC for Continuous Models," on page 26.

The overall scaled residual value (see Sections 8.6.2 and Section 8.7.2 for description of scaled residual derivation for continuous and dichotomous responses, respectively), and it corresponding overall p-value are indications of that "closeness". If the overall p-value is larger than some predetermined critical p-value, then the user may be able to conclude that the model is appropriate to model the data. The critical -value used by EPA is generally 0.1, but is sometimes relaxed to 0.05 for Multistage model when it is applied to cancer data (U.S. EPA, 2012).

8.3 Model Parameter Estimates

The model parameter estimates are provided in the Model Parameters table of the Model-Option worksheets of the Results Workbook. This table includes both the estimates for the true parameter values as well as their estimated standard errors. The standard errors are given for two reasons:

1. If standard errors are extraordinarily high, then the user may suspect that the probability function may not have reached a maximum, and they may want to use different starting points. There is not a guarantee if these are high that the function has not, in fact, been maximized. The user should use this in conjunction with other output to reach a decision.
2. To make inferences about the population parameters themselves. Under certain assumptions, the user may be able to formulate tests for the true value of the parameter.

8.4 Graphic Output from Models

The graphic output plot should display in the Summary and individual Model-Option worksheets of the Results Workbook along with the tabular results.

Figure 3. Results plot
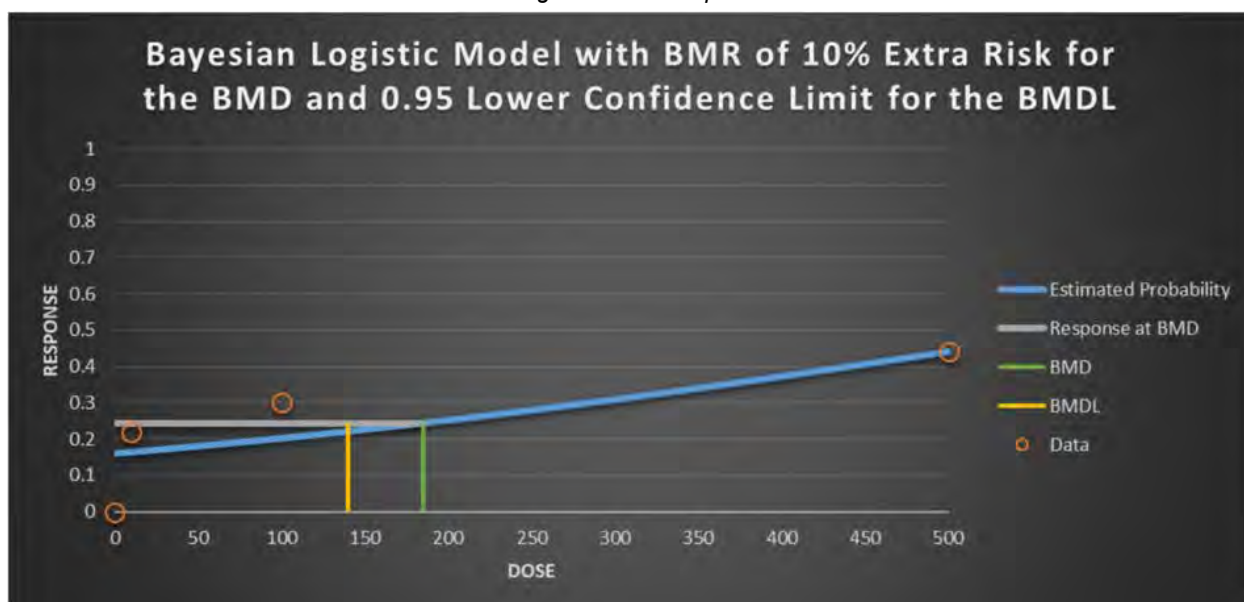
- The BMD and BMDL are indicated by the green and yellow vertical lines, respectively, and are associated with the user-selected benchmark response (BMR), the horizontal grey line.
- The BMD curve estimated by the model is represented by a blue line.
- Data points are shown as orange circles with their individual group confidence intervals (see the next section on error bar calculations for more information).
- The graphic display features can be modified using Excel edit features.

8.5 Plot Error Bar Calculations (Not in BMDS 3.0 Release)

Error bars will be added to BMDS 3.0 plots in the next BMDS update.

8.5.1 Continuous Models

BMDS uses a single error bar plotting routine for all continuous models.

1. The plotting routine calculates the standard error of the mean (SEM) for each group. The routine divides the group-specific observed variance (obs standard deviation squared) by the group-specific sample size.
2. The routine then multiplies the SEM by the Student-T percentiles (2.5th percentile or 97.5th percentile for the lower and upper bound, respectively) appropriate for the group-specific sample size (i.e., having degrees of freedom one less than that sample size). The routine adds the products to the observed means to define the lower and upper ends of the error bar.

8.5.2 Dichotomous Models

The error bars shown on the plots of dichotomous data are derived using a modification of the Wilson interval (based on the score statistic) but with a continuity correction method ([Fleiss et al., 2003](#)). For the upper bound, the calculation finds the proportion, pi, such that

$$\frac{|p - p_i| - \frac{1}{2n}}{\sqrt{\frac{p_i \times (1 - p_1)}{n}}} = z$$

where

- p is the observed proportion
- n is the total number in the group in question
- $z = Z_{1-\frac{\alpha}{2}}$ is the inverse standard normal cumulative distribution function evaluated at $1 - \frac{\alpha}{2}$

This leads to equations for the lower and upper bounds of:

- $LL = \dfrac{(2np + z^2 - 1) - z\sqrt{z^2 - (2 + \frac{1}{n}) + 4p(nq+1)}}{2(n+z^2)}$
- $UL = \dfrac{(2np + z^2 + 1) + z\sqrt{z^2 + (2 - \frac{1}{n}) + 4p(nq-1)}}{2(n+z^2)}$

where $q = 1 - p$.

The error bars shown in BMDS plots use alpha = 0.05 and so represent the 95% confidence intervals on the observed proportions (independent of model).

8.5.3 Nested Models

The error bars shown for the plots of nested data are calculated in the same way as those for dichotomous data. However, a Rao-Scott transformation is applied prior to the calculations to express the observations in terms of an "effective" number of affected divided by the total number in each group (the format required for the confidence intervals of simple dichotomous responses).

8.6 Outputs Specific to Continuous Models

8.6.1 Asymptotic Correlation Matrix of Parameter Estimates (Not in BMDS 3.0 Release)

This feature will be added to BMDS 3.0 in the next BMDS update.

This table in the individual model results provides the user with a matrix of correlation estimates between each of the parameters. Again, if these values seem to be high (in this case, very close to 1, in absolute value), there may have been a problem in the maximization. However, as stated before, high correlation does not confirm that the process of maximization did, in fact, fail.

Note The parameter standard errors and the correlation matrix elements are based on a variance-covariance (VCV) matrix obtained by inverting the negative of the Hessian matrix (the Fisher-observed information matrix). That matrix is made up of second partial derivatives of the log-likelihood, with respect to the model parameters. For all the continuous models, the partials are derived using a finite difference approximation to those derivatives.

8.6.2 Table of Data and Estimated Values of Interest

8.6.2.1 Goodness of Fit Table

This table in the individual model results gives a listing of the data as well as estimated means and standard deviations from the model. This is a good place for the user to look, along with the Tests of Fit and Maximum Likelihood below, to judge the appropriateness of the model. If a model fits well, the observed and estimated means should be relatively close. The scaled residual values printed in the final column of the table are defined as follows:

$$\frac{(\text{Obs.Mean} - \text{Predicted Mean})}{SE},$$

where the Predicted Mean is from the model and SE equals the estimated standard deviation (square root of the estimated variance) divided by the square root of the sample size.

The overall model should be called into question if the scaled residual value for any individual dose group, particularly the control group or a dose group close to the BMD estimate, is greater than 2 or less than -2.

8.6.2.2 Likelihoods of Interest Table

BMDS uses likelihood theory to estimate function parameters and ultimately to make inferences based on risk assessment data. Maximum likelihood is the process of estimating the model parameters; the likelihood function is as large as possible (maximized) given the form of the model under consideration and the data.

In other words, parameter values are "chosen" such that the subject model (e.g., polynomial or power) obtains the best possible fit to the data, given the constraints of the model's parameter structure.

For example, suppose one wishes to fit a second-degree polynomial model with a constant variance to a dataset. The form of this model would be:

$$Y = b_0 + b_1 * X + b_2 * X^2$$

The parameters we wish to estimate in this case would be b_0, b_1, and b_2 as well as the constant variance parameter, call it σ^2. To estimate these parameters, BMDS uses maximum likelihood procedures, the result being a vector of parameters that maximizes the likelihood function for the model specified.

The "Log(likelihood)" value given for BMDS modeling results is the maximum value of the natural logarithm of the likelihood function.

Also note that there are an associated number of parameters for each likelihood calculated. The number of parameters reported for the model under consideration is the total number possible for the model minus any parameter estimates that have values on the bounds set for their estimation (either bounds specified by the user or those inherent to the model).

In the example above, if all 4 parameters were estimated, and did not equal a bound (e.g., did not equal 0 for the b parameters), the number of parameters reported for the fitted model likelihood is 4.

The BMDS Results Workbook worksheets for continuous models provide five likelihood and AIC values that may be of interest to the user. These values are used in asymptotic Chi-Square tests of fit. Each of these likelihood values represents a model a user may consider in the analysis of the data. The five models are summarized in the following table.

Note	BMDS 3.0 handles Akaike Information Criterion (AIC) calculations somewhat differently from BMDS 2.x to facilitate comparing models with different likelihoods (i.e., normal vs. lognormal). For more details, refer to Section 7.3, "AIC for Continuous Models," on page 26.

Table 1. *Likelihood values and models*

Model	Description
A1: "Full" Constant Variance Model	$Y_{ij} = \mu_i + e_{ij},$ $Var\{e_{ij}\} = \sigma^2$
A2: "Fullest" Model	$Y_{ij} = \mu_i + e_{ij},$ $Var\{e_{ij}\} = \sigma_i{}^2$
A3: "Full" Model with variance structure specified by the user	$Y_{ij} = \mu_i + e_{ij},$ $Var\{e_{ij}\} = \alpha \times \mu_i{}^\rho$
R: "Reduced" Model	$Y_i = \mu + e_i,$ $Var\{e_i\} = \sigma^2$
Fitted Model	The user specified model

Model A1 estimates separate and independent means for the observed dose groups (it is "full" or "saturated" in that respect) but posits a constant variance over those groups.

Model A2 is the "fullest" model in that it estimates separate and independent means for the observed dose groups (as in Model A1) and it also estimates separate and independent variances for those groups. There is no assumed functional relationship among the means or among the variances across dose groups. This model is often referred to as the "saturated" model (it has as many mean and variance parameters as there are dose groups). The log-likelihood obtained for this model is the maximum attainable, for the data under consideration.

Model A3 is similar to model A2, and may only differ with respect to its variance parameters. Model A2 estimates separate and independent means for the observed dose groups (like A1). If the user specifies a constant variance for the fitted model, then model A3 will also assume that and it becomes identical to Model A1. If the user assumes a non-constant variance for the fitted model, then Model A3 will also assume the same functional form for the variance.

The reduced model (R) is the model that implies no difference in mean or variance over the dose levels. In other words, it posits a constant mean response level with the same variance around that mean at every dose level.

The last model, the fitted model, is the user-specified model (e.g., power or polynomial, among others). A user may have reason to believe that a certain model may describe the data well, and thus uses it to calculate the BMD and BMDL.

Refer to the next section for a description of how these models are used to test certain hypotheses about the data.

8.6.2.3 Tests of Interest Table

BMDS provides four different Tests of Fit that the user may use to determine an appropriate model for fitting their data. These Tests of Fit are based on asymptotic theories of the likelihood ratio.

Without getting too technical, the likelihood ratio is just the ratio of two likelihood values, many of which are given in the BMDS output. Statistical theory proves that $-2 * log(likelihood\ ratio)$ converges to a Chi-Square random variable as the sample size gets large and the number of dose levels gets large. These values can in turn be used to

obtain approximate probabilities to make inferences about model fit. Chi-Square tables can be found in almost any statistical book.

Each of the five models described in the previous section on likelihood has a likelihood value. The BMDS program uses these values to create ratios from two models that form a meaningful test. Suppose the user wishes to test two models, A and B, for fit. One assumption that is made for these tests is that model A is "nested within" Model B, i.e., that Model B can be simplified (via restriction of some parameters in Model B) in such a way that the simplified model is Model A. This implies that Model A has fewer varying parameters. As an example, consider that the linear model is a "simpler" or "nested" model relative to the power model because the linear model has the power parameter restricted to be equal to 1.

> *Note*   The model with a higher number of parameters is always in the denominator of this ratio.

Now, using the theory, $-2 \times \log\{\frac{L(A)}{L(B)}\}$ approaches a Chi-Square random variable. This can be simplified by using the fact that the log of a ratio is equal to the difference of the logs, or simply put,

$$-2 \times \log\left\{\frac{L(A)}{L(B)}\right\} = -2 \times (\log\{L(A)\} - \log\{L(B)\}) = 2 \times \log\{L(B)\} - 2 \times \log\{L(A)\}$$

The likelihood values given by BMDS are in fact the log-likelihoods, $log\{L(B)\}$ and $log\{L(A)\}$, so this likelihood ratio calculation becomes just a subtraction problem. This value can then in turn be compared to a Chi-Square random variable with a specified number of degrees of freedom.

As mentioned in the section on likelihood, each log likelihood value has an associated number of parameters. The number of degrees of freedom for the Chi-Square test statistic is merely the difference between the two model parameter counts. In the mini-example above, suppose Model A has 5 fitted parameters, and that Model B has 8. In this case, the Chi-Square value you would compare this to would be a Chi-Square with 8 - 5 = 3 degrees of freedom.

In the A vs B example, what is exactly being tested? In terms of hypotheses, it would be:

   H_0: A models the data as well as B

   H_1: B models the data better than A

Keeping these tests in mind, suppose $2 \times log\{L(B)\} - 2 \times log\{L(A)\} = 4.89$ based on 3 degrees of freedom. Also, suppose the rejection criteria is a Chi-Square probability of less than .05. Looking on a Chi-Square table, 4.89 has a p-value somewhere between .10 and .25. In this case, H_0 would not be rejected, and it would seem to be appropriate to model the data using Model A. BMDS automatically does the "table look-up" for the user, and provides the p-value associated with the calculated log-likelihood ratio having degrees of freedom as described above.

The BMDS software provides four default tests. BMDS provides interpretation of the test results, based on p-values that have been selected by EPA. However, the computed p-values are presented so that the user is free to use any rejection criteria they want. Each of the four default tests provided for any of the continuous models is discussed in some detail below.

Test 1 (A2 vs R): Tests the null hypothesis that responses and variances don't differ among dose levels. If this test fails to reject the null hypothesis, there may not be a dose-response.

This test compares Model R (the simpler model) to Model A2. Model R is a simpler A2 (or nested within A2) since R can be obtained from A2 by restricting all the mean parameters

to be equal to one another and restricting all the variance parameters to be equal to one another. If this test fails to reject the null hypothesis, then there may not be a dose-response, as the inference would be that the simpler model (R) is not much worse than the saturated model. The default p-value for the test (as reported in the Tests of Interest section of the output) is 0.05. A p-value less than 0.05 is an indication that there is a difference between response and/or variances among the dose levels and supports a conclusion to model the data. A p-value greater than 0.05 is an indication that the data may not be suitable for dose-response modeling.

Test 2 (A1 vs A2): Tests the null hypothesis that variances are homogeneous. If this test fails to reject the null hypothesis, the simpler constant variance model may be appropriate.

This test compares A1 (the simpler model) to Model A2. Model A1 is a simpler A2 (or nested within A2) since A1 can be obtained from A2 by restricting all the variance parameters to be equal to one another. If this test rejects the null hypothesis, the inference is that the constant variance assumption is incorrect and a modeled variance is necessary to adequately represent the data. The default p-value for the test (as reported in the Tests of Interest section of the output) is 0.05. A p-value less than 0.05 is an indication that the user should consider running a non-homogeneous variance model. A p-value greater than 0.05 is an indication that a constant variance assumption may be suitable for the dose-response modeling.

Test 3 (A3 vs A2): Tests the null hypothesis that the variances are adequately modeled. If this test fails to reject the null hypothesis, it may be inferred that the variances have been modeled appropriately.

Here, the test is one to see if the user-specified variance model, is appropriate. If the user-specified variance model is "constant variance," then Models A1 and A3 are identical; this test is the same as Test 2, with the same interpretation. If the user-specified variance model is nonconstant ($\sigma_i^2 = \alpha \times \mu_i^\rho$), this test determines if that equation appears adequate to describe the variance across dose groups. Model A3 is the simpler version of Model A2 obtained by constraining the variances to fit the nonconstant variance equation. The default p-value for the test (as reported in the Tests of Interest section of the output) is 0.05. A p-value less than 0.05 is an indication that the user may want to consider a different variance model. A p-value greater than 0.05 supports the use of modeled variance for the dose-response modeling.

Test 4 (Fitted vs A3): Tests the null hypothesis that the model for the mean fits the data. If this test fails to reject the null hypothesis, the user has support for the selected model.

This test compares the Fitted Model to Model A3. The Fitted Model is as simpler Model A3 (or nested within Model A3) because it can be obtained by restricting the means (unrestricted in A3) to be described by the dose-response function under consideration. If this test fails to reject the null hypothesis, the inference is that the fitted model is adequate to describe the dose-related changes in the means (conditional on the form of the variance model; the form of the variance model is the same for the Fitted Model and Model A3). Failure to reject the null hypothesis is associated with the inference that the restriction of the means to the shape of the dose-response function under consideration is adequate. The default p-value for the test (as reported in the Tests of Interest section of the output) is 0.1. A p-value less than 0.1 is an indication that the user may want to try a different model (i.e., the fit of the Fitted Model is not good enough). A p-value greater than 0.1 is an indication that the Fitted Model appears to be suitable for dose-response modeling.

8.7   Outputs Specific to Dichotomous Models

8.7.1 Asymptotic Correlation Matrix of Parameter Estimates (Not in BMDS 3.0 Release)

This feature will be added to BMDS 3.0 in the next BMDS update.

This table in the individual model results provides the user with a matrix of correlation estimates between each of the parameters. Again, if these values seem to be high (in this case, very close to 1), there may have been a problem in the maximization. Also, as stated before, high correlation does not confirm that the problem of maximization in fact failed. The Weibull model, for instance, tends to give high correlation between the slope and power parameters, even when the likelihood was maximized.

Note   The parameter standard errors and the correlation matrix elements are based on a variance-covariance (VCV) matrix obtained by inverting the negative of the Hessian matrix (the Fisher-observed information matrix). That matrix is made up of second partial derivatives of the log-likelihood, with respect to the model parameters.

For all the dichotomous models, except for the multistage model, the partials are derived using a finite difference approximation to those derivatives.

For the multistage model, the partial derivatives are computed analytically (i.e., without approximating their values through the finite-difference method).

8.7.2 Goodness of Fit

This table in the individual model results gives both a listing of the data, model response estimates and scaled residuals. This is a good place for the user to look outside of the overall goodness-of-fit statistics reported in the Results Workbook worksheet's Benchmark Dose Table (Section 8.2) and the Analysis of Deviance table (Section 8.7.3) to judge the appropriateness of the model. The table lists estimated probabilities, the expected and observed number of affected animals and scaled residuals for each dose group. If a model fits well, the observed and expected number of affected animals should be relatively close.

The scaled residual values printed at the end of the table are defined as follows:

$$\frac{(\text{Obs} - \text{Expected})}{SE},$$

where "Expected" is the predicted number of responders from the model and SE equals the estimated standard error of that predicted number. For these models, the estimated standard error is equal to $\sqrt{[n \times p \times (1 - p)]}$, where

n is the sample (litter) size, and

p is the model-predicted probability of response.

The overall model should be called into question if the scaled residual value for any dose group, particularly the control or low dose group, is greater than 2 or less than -2.

8.7.3 Analysis of Deviance Table

The analysis of deviance table lists three maximum likelihood values.

- The first is the "full model". The full model would be any model that would perfectly fit all the positive response proportions at the dose levels specified by the user.
- The second model is the "fitted model" maximum likelihood value. This is the value of the maximum likelihood function for the selected model and using the estimated parameter values.
- The last likelihood value is the "reduced model" value, which would be the value of the likelihood function if all data points where assumed to come from the same population with the same population parameter. That is, for each dose level, the actual probability of an adverse effect would be the same. These values are just the likelihood functions evaluated according to the assumptions made at each step (i.e., the model assumption for the fitted model).

Next to the likelihood values there are three values: Deviance, degrees of freedom (DF), and P-value.

- The Deviance is the difference between the fitted or reduced model and the full model likelihood values. This deviance measures if the smaller model (i.e., the fitted or reduced model) describes the data as well as the full model does. This deviance is then used to formulate a Chi-Square random variable that tests exactly that. The user may choose a rejection level (0.05 is common) to test if the model fit is appropriate.
- The p-value for testing if the fitted model adequately describes the data is given next to the fitted model likelihood, and the user can reject or not reject a hypothesis according to the p-value given.
- The reduced model p-value would be used in the same way, but here the user would be testing if there is in fact a dose/response relationship where the true population proportion is a function of dose, as opposed to a single population with one parameter (the proportion of affected animals).

It will often happen that several models provide an adequate fit to a given dataset. These models may be essentially unrelated to each other (for example a logistic model and a probit model often do about as well at fitting dichotomous data) or they may be related to each other in the sense that they are members of the same family that differ in which parameters are fixed at some default value. One can consider the log-logistic, the log-logistic with non-zero background, and the log-logistic with threshold and non-zero background to all be members of the same family of models. Generally, within a family of models, as additional parameters are introduced the fit will appear to improve. Goodness-of fit statistics presented in the main body of the Analysis of Deviance Table can be used to compare such related models, but are not designed to compare unrelated models. Alternative approaches are need for selecting between models that are not related (not in the same family).

8.7.4 Cancer Slope Factor—Restricted Multistage Model Only

Some additional assessment tools are imparted by the restricted Multistage model for use with cancer response data. The output page for the restricted Multistage model includes an estimate of the cancer slope factor (CSF), defined by EPA as the linear slope between the extra risk at the BMDL(10) and the extra risk at background (generally 0 dose). The restricted Multistage model plot also includes a dashed line representing this linear slope.

8.8 Outputs Specific to Bayesian Dichotomous Models

To compare the difference between any two Bayesian models, the unnormalized Log Posterior Probability (LPP) is given, which allows the computation of Bayes factors to compare any two models. To compare any two models, one takes the difference between the two LPP and exponentiates this difference. For example, if one wishes to compare the Log-Logistic (Model A) model (i.e., LPP_A) to the multistage 2nd degree model (i.e., LPP_B) one estimates the Bayes factor as

$$BF = \exp(LPP_A - LPP_B),$$

which assumes that both models have equal probability *a priori.* This value is then interpreted as the posterior odds one model is more correct than the other model, and is used in Bayesian hypothesis testing. In the example above, if the Bayes Factor was 2.5, the interpretation would be that the Log-logistic model is *a posteriori* 2.5 times more likely multistage model. When these values are normalized into proper probabilities, they are equivalent the posterior model probabilities given in model averaging except they assume equal prior probability *a priori*. The table below, adapted from Jeffreys (Jeffreys, 1998) is a common interpretation of Bayes Factors.

Table 2. *Bayes Factors*

Bayes Factor	Strength of evidence for H_A
< 1	negative (supports HB)
1 to 3.2	not worth mentioning
3.2 to 10	substantial
10 to 31.6	strong
31.6 to 100	very strong
100	decisive

For BMDS 3.0, all LPP and corresponding posterior model probabilities are computed using the Laplace approximation. This value is different from the commonly used Bayesian Information Criterion (BIC), and the two should not be confused based upon other model averaging approaches, which use the BIC exclusively. Posterior probabilities formed from the BIC are $O(1)$ estimators, where posterior probabilities formed using the Laplace approximation are $O(n^{-1})$. This means the latter approximation goes to the true posterior model probability with increasing data and the former, using the BIC, may not go to the true value.

8.9 Outputs Specific to Nested Models

8.9.1 Analysis of Deviance Table

The analysis of deviance table in the individual model results lists three maximum likelihood values.

- The first is the "full model". The full model would be any model that would perfectly fit all the positive response proportions at the dose levels specified by the user.
- The second model is the "fitted model" maximum likelihood value. This is the value of the maximum likelihood function for the selected model and using the estimated parameter values.

- The last likelihood value is the "reduced model" value, which would be the value of the likelihood function if all data points where assumed to come from the same population with the same population parameter. That is, for each dose level, the actual probability of an adverse effect would be the same.

These values are just the likelihood functions evaluated according to the assumptions made at each step (i.e., the model assumption for the fitted model).

Next to the likelihood values there are three values: Deviance, degrees of freedom (DF), and P-value. These are asymptotic Chi-Square tests that investigate the appropriateness of the model fit, as well as the reduced model.

- The Deviance is the difference between the fitted or reduced model and the full model likelihood values. This deviance measures if the smaller model (i.e., the fitted or reduced model) describe the data as well as the full model does. This deviance is then used to formulate a Chi-Square random variable that tests exactly that. The user may choose a rejection level (0.05 is common) to test if the model fit is appropriate.
- The p-value for testing if the fitted model adequately describes the data is given next to the fitted model likelihood, and the user can reject or not reject a hypothesis according to the p-value given.
- The reduced model p-value would be used in the same way, but here the user would be testing if there is in fact a dose/response relationship where the true population proportion is a function of dose, as opposed to a single population with one parameter (the proportion of affected animals).

8.9.2 Goodness of Fit Information—Litter Data and Grouped Data

Both tables provide a listing of the data, expected and observed responses and scaled residuals (observed - expected).

The "Litter Data" table contains this information for each litter.

To obtain the "Group Data" table, the Litter Data were sorted on Dose (first), and by Litter Specific Covariate within Dose. Within dose, litters adjacent to each other with respect to Litter Specific Covariate were grouped together until the expected number of affected pups was at least one. This grouping was done prior to the estimation of an overall Chi-Square and p-value to improve the validity of the Chi-Square approximation for the goodness of fit statistic. Goodness of Fit statistics. Both tables list estimated probabilities, the expected and observed number of affected animals and scaled residuals for each dose group. If a model fits well, the observed and expected number of affected animals should be relatively close. The overall Chi-Square value, and it corresponding p-value are an indication of that "closeness". If the p-value is larger than some predetermined critical p-value, then the user may be able to conclude that the model is appropriate to model the data.

The scaled residual values printed at the end of the table are defined as follows:

$$\frac{(\text{Obs} - \text{Expected})}{SE},$$

where "Expected" is the predicted number of responders from the model and SE equals the estimated standard error of that predicted number. For these models, the estimated standard error is equal to $\sqrt{[n \times p \times (1 - p) \times (\theta \times (n - 1) + 1)]}$,

- n is the sample (litter) size,
- p is the model-predicted probability of response, and
- θ is the model-predicted intra-litter correlation coefficient.

The overall model should be called into question if the scaled residual value for any individual dose and litter-specific covariate combination, particularly for the control or a low dose group, is greater than 2 or less than -2.

9.0 Model Descriptions

9.1 Models Included in BMDS 3.0

The following tables detail the models and version numbers contained in this BMDS version.

Frequentist Dichotomous Models	Version (Date)
Dichotomous Hill	1.0 (09/30/2018)
Gamma Model	1.0 (09/30/2018)
Logistic and Log-Logistic Models	1.0 (09/30/2018)
Probit and Log Probit Models	1.0 (09/30/2018)
Multistage Model	1.0 (09/30/2018)
Weibull and Quantal Linear Models	1.0 (09/30/2018)

Bayesian Dichotomous Models	Version (Date)
Dichotomous Hill	1.0 (09/30/2018)
Gamma Model	1.0 (09/30/2018)
Logistic and Log-Logistic Models	1.0 (09/30/2018)
Probit and Log Probit Models	1.0 (09/30/2018)
Multistage Model	1.0 (09/30/2018)
Weibull and Quantal Linear Models	1.0 (09/30/2018)

Frequentist Continuous Models	Version (Date)
Exponential Model	1.0 (09/30/2018)
Hill Model	1.0 (09/30/2018)
Polynomial and Linear Models	1.0 (09/30/2018)
Power Model	1.0 (09/30/2018)

Nested Dichotomous Models	Version (Date)
NCTR Model	2.13 (04/27/2015)
Nlogistic Model	2.20 (04/27/2015)

Bayesian Model Averaging	Version (Date)
Bayesian Model Averaging	1.0 (09/30/2018)

Multi-Tumor (MS_Combo) Model	Version (Date)
Multiple tumor analysis; combining restricted Multistage (cancer model) model runs over different tumors	1.8 (04/30/2014)

9.2   Model Types and Abbreviations

BMDS uses the following naming conventions for model abbreviations.

Continuous

Exponential	exp
Hill	hil
Linear	lin
Polynomial	ply
Power	pow

Dichotomous

Gamma	gam
Logistic	log
LogLogistic	lnl
LogProbit	lnp
Multistage	mst
Probit	pro
Weibull	wei
Quantal Linear	qln
Dichotomous Hill	dhl

Nested

NCTR	nctr
Nested Logistic	nln

Bayesian Model Averaging

bma

Multi-tumor

    MS_Combo   multi

9.3   Optimization Algorithms Used in BMDS

The NLopt optimization library is used for BMDS 3.0.

Several optimization algorithms available in the library are used to ensure reliability of the estimation:

- For global optimization involving the maximum likelihood or maximum a-posteriori estimation, the L-BFGS method is attempted first. If it fails to converge, gradient free algorithms "subplex" and "BOBYQA" algorithms are then attempted.
- For profiling, when only non-linear inequality constraints are needed, the COBYLA and MMA approaches are used and compared. In the case the methods return different optimum, the values producing the larger of the two is used.
- For equality-constrained optimization, the augmented Lagrangian algorithm is used and either the L-BFGS, BOBYQA, or the "subplex" algorithm is used in the local optimization step. When two approaches produce different results, the values producing the larger optimum is used.

NLopt 2.4.1 was used when developing the BMDS 3.0 code. This version is available for download from the NLopt GitHub site.

For more information regarding the algorithms, refer to the NLopt documentation site.

9.4   Frequentist Continuous Model Descriptions

9.4.1 Special Considerations for Models for Continuous Endpoints in Simple Designs

Models in this section are for continuous endpoints, such as weight or enzyme activity measures, in simple experimental designs that do not involve nesting or other complications. The models predict the mean value of the response, μ (dose), expected for a given dose.

Models for continuous endpoints require consideration of more details than do those for dichotomous endpoints in similar designs. While for dichotomous models, we normally model the incidence of adversely affected individuals, and so expect the response to increase with increasing dose, in continuous models the change in a measure is modeled without regard for "adversity," and the response may increase or decrease. Thus, just what constitutes an adverse change, and how to specify it, must be made explicit. The models in BMDS allow that specification to be made in several ways, which will be described below (BMD Computation).

Another important contrast with dichotomous models is the nature of the probability distribution of response. In dichotomous models, the nature of the experimental design guarantees that the binomial probability distribution is appropriate. There are many more options for continuous distributions, however. In the current version of BMDS, the distribution of continuous measures is assumed to be normal, except for the Exponential Models, for which the user may assume either a normal or a lognormal distribution (see Section 9.4.5, "Lognormal Distributions," on page 45 for more information). Moreover, for all models and normally distributed data, one may assume either a constant variance

(that is, the variance is the same regardless of dose group), or a variance that changes as a power function of the mean value:

$$\sigma_i{}^2 = \alpha[\mu(dose_i)]^\rho,$$

which is the modeled variance for the *ith* dose group. the expression μ (*dose$_i$*) is the observed mean (from the model) for the *ith* dose group, and α and ρ are estimated parameters. This formulation allows for several commonly encountered situations. For example, if $\rho = 2$, then the coefficient of variation is constant, a common situation especially for biochemical measures; if $\rho = 1$, then the variance is proportional to the mean, which is sometimes appropriate for large counts (especially if the constant of proportionality, *k*, is 1.0). When a lognormal distribution is assumed, the Exponential Models assume a constant (log-scale) variance, equivalent to a constant coefficient of variation.

9.4.2 Likelihood Function

Suppose there are *g* doses,

$$dose_1, \dots, dose_g$$

with N_i subjects per dose group, and that Y_{ij} is the measurement for the *jth* subject in the *ith* dose group. The form of the log-likelihood function depends upon whether the variance is assumed to be constant, or to vary among doses.

For constant variance, the log-likelihood function is:

$$L = -\frac{g}{2}\ln(2\pi) - \sum_{i=1}^{g}\left[\frac{N_i}{2}\ln\sigma_i{}^2 + \frac{(N_i-1)s_i{}^2}{2\sigma_i{}^2} + \frac{N_i\left(\bar{y}_i - \mu(dose_i)\right)^2}{2\sigma_i{}^2}\right]$$

Where

$$\sigma_i^2 = \frac{\sum_{j=1}^{N_i}(y_{ij} - \bar{y}_i)^2}{N_i - 1}$$

is the sample variance for the *ith* dose group,

$$y_i = \frac{\sum_{j=1}^{N_i} y_{ij}}{N_i}$$

is the sample mean for the *ith* dose group, *g* is the number of doses, *N$_i$* is the number of subjects in the *ith* dose group, and σ² the variance which is same in all dose groups. Generally, σ² and the parameters hidden here in λ( ) are to be estimated.

If the variance is allowed to be a power function of the mean, the log-likelihood function is:

$$L = \sum_{i=1}^{g}\frac{N_i}{2}\ln\alpha + \frac{N_i\rho}{2}\ln[\mu(x_i)] + H_i]$$

Where

$$H_i = \frac{A_i}{2\alpha[\mu(dose_i)]^\rho} - \frac{B_i}{\alpha[\mu(dose_i)]^{\rho-1}} + \frac{N_i}{2\alpha[\mu(dose_i)]^{\rho-2}}$$

With

$$A_i = (N_i)s_i^2 + N_i\bar{y}_i^2$$
$$B_i = N_i\bar{y}_i$$

The upper bound for the power parameter in the Hill and power models has been (somewhat arbitrarily) set to 18. That value was selected because it represents a very high degree of curvature that should accommodate almost every dataset, even ones with very (or absolutely) flat dose-response at low doses followed by a very steep dose-response at higher doses.

If the power parameter for the Hill or power model is reported equal to 18 and the warning "*... hit a bound ...* " appears, the parameter estimates are maximum likelihood estimates only in the restricted sense that the power parameter has been assigned a value and the other parameters are MLEs conditional on that assigned value. Such model results are not strictly comparable with others in terms of AIC. In such a case, the BMD and BMDL could depend on the choice of power parameter; thus, sensitivity analysis is indicated if one intends to rely on the reported BMD or BMDL.

Note BMDS 3.0 handles Akaike Information Criterion (AIC) calculations somewhat differently from BMDS 2.x to facilitate comparing models with different likelihoods (i.e., normal vs. lognormal). For more details, refer to Section 7.3, "AIC for Continuous Models," on page 26.

9.4.3 BMD Computation

In the continuous models, the benchmark dose is always the dose that results in a prespecified change in the mean response. The change can be expressed in several ways:

- an absolute change in the mean (Abs. Dev.);
- a change in the mean equal to a specified number of control standard deviations (Std. Dev);
- a specified fraction of the control group mean (Rel. Dev.);
- a specified value for the mean at the BMD (i.e., not a change, but a fixed value) (Point);
- a change equal to a specified probability of increase above background. (Hybrid) [A probabilistic definition equivalent to extra risk in the dichotomous model suite].

Symbolically, these are (where δ represents the BMRF designated by the user):

$$|\mu(BMD) - \mu(0)| = \begin{cases} \delta, Abs.Dev \\ \delta \cdot \hat{\sigma}_1, Std.Dev \\ \delta \cdot \mu(0), Rel.Dev \end{cases}$$

$$\mu(BMD) = \delta \; Point$$

$$\frac{\Pr(X > X_0|BMD) - \Pr(X > X_0|0)}{1 - \Pr(X > X_0|0)} = \delta\ Hybrid$$

9.4.4 BMDL and BMDU Computation

The general approach to computing the confidence limits for the BMD (called the BMDL and BMDU here) is the same for all the models in BMDS, and is based on the asymptotic distribution of the likelihood ratio (Crump and Howe, 1985). Two different approaches are followed in these models. In one (used for the power model), the equations that define the benchmark response in terms of the benchmark dose and the dose-response model are solved for one of the model parameters. The resulting expression is substituted back into the model equations, with the effect of re-parameterizing the model so that BMD appears explicitly as a parameter. A value for BMD is then found such that, when the remaining parameters are varied to maximize the likelihood, the resulting log-likelihood is less than that at the maximum likelihood estimates by exactly

$$\frac{\chi^2_{1,1-2\alpha}}{2}$$

In the polynomial, Hill, and exponential models, it is impractical or impossible to explicitly reparameterize the dose-response model function to allow BMD to appear as an explicit parameter. For this model, the BMR equation is used as a non-linear constraint, and the minimum value of BMD is determined such that the log-likelihood is equal to the log-likelihood at the maximum likelihood estimates less

$$\frac{\chi^2_{1,1-2\alpha}}{2}$$

Occasionally, the following error message may appear for a model: "BMDL computation is at best imprecise for these data." This is a flag that convergence for the BMDL was not "successful" in the sense that the required level of convergence (< 1e-3 relative change in the target function by the time the optimizer terminates) has not been achieved.

9.4.5 Lognormal Distributions

In previous versions of BMDS, continuous data were always assumed to be normally distributed. In the current version of BMDS, for the exponential models only, the user has the option of specifying that the continuous data being analyzed are lognormally distributed. Lognormal distributions are appropriate only for data that are strictly positive and may be preferable for such data (since the normal distribution allows, in theory, both positive and negative values, no matter what the mean and standard deviation). When a lognormal distribution is specified, the models assume a constant log-scale variance, which is equivalent to an assumption of a constant coefficient of variation (CV).

The likelihood function shown above is then correct for data on the log scale (log-transformed) and is the basis for fitting the log-transformed version of the model in question. That is, if $\mu L(dose)$ is the log-scale mean as a function of dose, the model being fit is $\mu L(dose) = ln\{m(dose)\}$, where m(dose) is the specified model (e.g., one of the exponential models parameterized as shown in the section on Exponential Models). Therefore, m(dose) will then be a description of the change in the median response as a function of dose since the anti-log of the log-scale mean is the median.

When the input data are summarized in terms of the sample mean and sample standard deviation (or standard error or variance), the exact likelihood of the data cannot be

determined if the data are lognormally distributed. In such cases, BMDS gives an approximate MLE solution by estimating the log-scale sample mean and log-scale sample standard deviation for each dose group as follows:

estimated log-scale sample standard deviation (sL): $sL = \sqrt{\left(\ln\left[1 + \frac{s^2}{m^2}\right]\right)}$

estimated log-scale sample mean (mL): $mL = \ln(m) - \frac{sL^2}{2}$

where m and s are the reported sample mean and sample standard deviation. When individual responses are available, the user may input those values (where the input data file will have two columns reporting the dose and the response for each experimental unit) and may request that the exact MLE solution be obtained (which the software does by first log-transforming the individual responses) or that the approximate solution using the estimates shown above be obtained (which the software does by first computing sample means and sample standard deviations). This option allows the user to compare estimates and determine the impact of the approximation or to provide consistency across datasets if some datasets have individual responses while others do not.

9.4.6 Continuous Frequentist Models

Table 3. The individual continuous models used and their respective parameters

Model	Parameters	Notes
Linear and Polynomial models $\mu(dose) = \beta_0 + \beta_1 dose + \beta_2 dose^2 + \cdots + \beta_n dose^n$ n is the degree of the polynomial, specified by user and must be a positive integer (maximum value = 21) The linear model is a special case of the polynomial model with n = 1	Alpha is α from variance model Rho is ρ from variance model $\beta_0 \ldots \beta_n$ are the polynomial coefficients	User can restrict the value of the polynomial coefficients. Restricting them to be either "non-positive" or "non-negative" guarantees that the resulting function will be strictly decreasing, strictly increasing, or perfectly flat (when all the coefficients are zero). If the coefficients are unrestricted (i.e., an unrestricted form of the model is run), more complicated shapes are possible, and, particularly as the degree of the polynomial approaches the number of dose groups minus one, the polynomial will often be quite "wavy".
Power $\mu(dose) = \gamma + \beta \times (dose)^\delta$	Alpha is α from variance model Rho is ρ from variance mode γ = control (intercept) β = slope δ = power	The Power parameter must be a positive number ≤ 18. If Power is restricted, the number must be > 1.) If $\delta < 1$, then the slope of the dose-response curve becomes infinite at the control dose. This is biologically unrealistic, and can lead to numerical problems when computing confidence limits, so several authors have recommended restricting $\delta \geq 1$.
Hill[1] $\mu(dose) = \gamma = \frac{v \times d^n}{K^n + d^n}$	γ = control (intercept) k = dose with half-maximal change n = power v = maximum change	k must be a positive number n must be a positive number ≤ 18. If n is restricted, the number must be > 1
Exponential[1,2] $M2, \mu(dose) = a \times e^{-b \times dose}$ $M3, \mu(dose) = a \times e^{-(b \times dose)^d}$ $M4, \mu(dose) = a \times (c - (c-1) \times e^{-b \times dose}$ $M5, \mu(dose) = a \times (c - (c-1) \times e^{-(b \times dose)^d}$	a = background response (>0) b = slope (>0) c = asymptote term: (>1 for increasing data, 0< c < 1 for decreasing data) d = power (>1)	The parameter "sign" is the indicator of the direction of change: +1 for data trending up, -1 for data trending down. It is very important that the user correctly specify the direction of change in the data - for the Exponential Models the "automatic" choice of adverse direction has not been included.

[1] BMDL estimates from models that have an asymptote parameter (including the Hill model) can be unstable when a wide range of parameter values can give nearly identical likelihoods. One indicator of that problem is that the estimated asymptotic response is far outside the range of the observed responses. The user should consult a statistician if this behavior is seen or suspected.

[2] RIVM (National Institute for Public Health and the Environment (Netherlands)). (2018). PROAST. Retrieved from https://www.rivm.nl/en/Documents_and_publications/Scientific/Models/PROAST

9.5   Frequentist Dichotomous Model Descriptions

9.5.1 Special Considerations for Models for Dichotomous Endpoints in Simple Designs

BMDS includes in this category models for dichotomous endpoints in which the observations are independent of each other. In these models, the dose-response model provides the probability that an animal will have an adverse response at a given dose. The actual number of animals that have an adverse response is assumed to be binomially distributed.

An example of such a dataset is a study in which adult animals are exposed to different concentrations of a toxicant and then evaluated for the presence of liver toxicity. For models for dichotomous endpoints in which the responses are nested (for example, pups in litters, and litters nested within doses), see the section on Nested Model Descriptions.

BMDS contains nine models for dichotomous endpoints (the Dichotomous Hill, Probit, Log-Probit, Logistic, Log-Logistic, Weibull, Quantal Linear, Gamma, and Multistage models). They may all be written in the form:

$$Prob\{response\} = \gamma + (1 - \gamma)F(dose;\ \alpha, \beta, \dots)$$

Here $F(dose;\ \alpha, \beta, \dots)$ is a cumulative distribution function and γ, α, β, . . ., are parameters to be estimated using maximum likelihood methods. Sometimes $Prob\{response\}$ is written as $P[dose; \gamma, \alpha, \beta, \dots]$ to indicate the relationship between the response probability and the dose as well as parameters. When the function $F(dose;\ \alpha, \beta, \dots)$ approaches zero as dose approach zero, the parameter γ represents the background incidence. In the Logistic and Probit models, *F*(*0*) is not zero, unlike in the Log-Logistic and Log-Probit models. In these models, γ is set to 0.

9.5.2 Special Options for Models

In addition to the options that are available to all dichotomous models, there may be model-specific options. Generally, these are options to restrict the legal range of a parameter or set of parameters. The range of a parameter may be restricted for two reasons:

- The slope of the dose-response curve becomes infinite at a dose of 0 if the parameter falls below 1, so that the default is to restrain that parameter to be at least 1, or
- The quantal polynomial dose-response curve can become non-monotonic if the coefficients are allowed to be negative, often resulting in the curve looking ''wavy'', so the default is to restrict the coefficients to be non-negative.

The applicable special options are listed in the sections for the specific models.

9.5.3 Likelihood Function

All models in the current version of BMDS are fit using maximum likelihood methods. This section describes the likelihood function used to fit the dichotomous models.

Suppose we employ k doses:

$$dose_1, dose_2, \dots, dose_k$$

and the total number of s and number of responding s in each dose group are

$$N_1, N_2, \ldots, N_k$$

and, respectively,

$$n_1, n_2, \ldots, n_k$$

The distribution of ni is assumed to be binomial with probability

$$p_i = p(dose_i;\ \theta), i = 1,2, \ldots k$$

where θ is a vector of parameters. Then the log-likelihood function L can be written as

$$L = \sum_{i=1}^{k} L_i(N_i, n_i, dose_i;\ \theta)$$

Where

$$L_i = n_i \ln(p_i) + (N_i - n_i) \ln(1 - p_i), i = 1,2, \ldots, k$$

The upper bound for the power parameter in the gamma and Weibull models, and the slope parameter for the log-probit and log-logistic models, has been (somewhat arbitrarily) set to 18. That value was selected because it represents a very high degree of curvature that should accommodate almost every dataset, even ones with very (or absolutely) flat dose-response at low doses followed by a very steep dose-response at higher doses.

If the power parameter for the gamma or Weibull model, or the slope parameter for the log-probit or log-logistic model, is reported equal to 18 and the warning "... hit a bound ... " appears, the parameter estimates are maximum likelihood estimates only in the restricted sense that the power parameter has been assigned a value and the other parameters are MLEs conditional on that assigned value. Such model results are not strictly comparable with others in terms of AIC. In such a case, the BMD and BMDL could depend on the choice of power parameter; thus, sensitivity analysis is indicated if one intends to rely on the reported BMD or BMDL.

9.5.4 BMD Computation

The BMD is computed as a function of the parameters of the model, which must have already been estimated. The BMDs for dichotomous models are expressed as the dose that would give an (estimated) increase in incidence of x% above the control incidence (where x is often in the range of 1 to 10 percent). This increase in incidence is referred to here as the "BMRF", for benchmark response factor. Note that, although the word "response" is used here, we are really talking about an increase of the incidence over the control incidence (added risk). The actual response associated with the BMR, will only be the same as the BMR when $P(0) = 0$. This is because to obtain the actual response associated with the BMR one must solve for $P(d)$ in the equation for added or extra risk.

Two formulations for computing the excess over background are in common use, the extra risk model and the additional risk model. In the extra risk model,

$$BMR = \frac{p(BMD; \gamma, \alpha, \beta, \ldots) - p(0; \gamma, \alpha, \beta, \ldots)}{(1 - p(0; \gamma, \alpha, \beta, \ldots))}$$

while in the additional risk model,

$$BMR = p(BMD; \gamma, \alpha, \beta, \dots) - p(0; \gamma, \alpha, \beta, \dots)$$

The equation appropriate to the risk type formulation that the user requests is solved to get the BMD for a specific model and dataset. Details of this computation are included in the descriptions of the models.

9.5.5 BMDL Computation

BMDS currently calculates one-sided confidence intervals, in accordance with current BMD practice. The general approach to computing the confidence limits for the BMD (called the BMDL and BMDU here) is the same for all the models in BMDS, and is based on the asymptotic distribution of the likelihood ratio (Crump and Howe, 1985) . Two different approaches are followed in these models. In one, the equations that define the benchmark response in terms of the benchmark dose and the dose-response model are solved for one of the model parameters. The resulting expression is substituted back into the model equations, with the effect of reparameterizing the model so that BMD appears explicitly as a parameter. A value for BMD is then found such that, when the remaining parameters are varied to maximize the likelihood, the resulting log-likelihood is less than that at the maximum likelihood estimates by exactly

$$\frac{\chi^2_{1,1-2\alpha}}{2}$$

In a few models, it is impractical or impossible to explicitly reparameterize the dose-response model function to allow BMD to appear as an explicit parameter. For these models, the BMR equation is used as a non-linear constraint, and the minimum value of BMD is determined such that the log-likelihood is equal to the log-likelihood at the maximum likelihood estimates less

$$\frac{\chi^2_{1,1-2\alpha}}{2}$$

9.5.6 Dichotomous Frequentist Models

Table 4. *The individual dichotomous models used and their respective parameters*

Model	Parameters	Notes
Multistage $$p(d) = \gamma + (1-\gamma)(1 - \exp[-\sum_{j=1}^{n} \beta_j dose^j])$$	γ = background (>0, <1) β = dose coefficients	Can restrict all β coefficients to be > 1. Doing so with guarantee that the dose-response function will either be perfectly flat or always increasing with no negative slopes. The degree of the multistage model is n. The maximum degree polynomial to fit is 23. Per EPA cancer guidance, when the Multistage model is used for cancer analyses (e.g., in the BMDS Multi-tumor model) all β coefficients are restricted to be positive.
Weibull $$p(d) = \gamma + (1-\gamma)(1 - \exp[-\beta d^\alpha])$$	γ = background (>0, <1) α = power (<18, ≥1 if restricted) β = slope (≥1 if restricted)	Restrict α ≥ 1. If α < 1, then the slope of the dose-response curve becomes infinite at the control dose. This is biologically unrealistic, and can lead to numerical problems when computing confidence limits, so several authors have recommended restricting α ≥ 1. The Quantal Linear model results from setting α = 1.
Gamma $$p(d) = \gamma + \frac{1-\gamma}{\Gamma(\alpha)} \int_0^{\beta d} t^{\alpha-1} \exp(-t)\, dt$$	γ = background (>0, <1) α = power (<18, ≥1 if restricted) β = slope (≥1 if restricted)	If α < 1, then the slope of the dose-response curve becomes infinite at the control dose. This is biologically unrealistic, and can lead to numerical problems when computing confidence limits, so several authors have recommended restricting α ≥ 1. Note for the unrestricted Gamma model the α > 0.2 for numerical reasons.
Logistic $$p(d) = \frac{1}{1+\exp[-\alpha-\beta d]}$$	α = intercept β = slope (>1 if restricted)	None
Log-Logistic $$p(d) = \gamma + \frac{1-\gamma}{1+\exp[-\alpha-\beta \log(d)]}$$	γ = background (>0, <1) α = power (<18, ≥1 if restricted) β = slope (>1 if restricted)	If the slope is allowed to be less than 1, the slope of the dose-response curve is infinite at zero dose.
Probit $$p(d) = \Phi(\alpha + \beta d), \text{ where } \phi(x) = \int_{-\infty}^{x} \phi(t)dt \text{ and } \phi(t) = \frac{1}{\sqrt{2\pi}} e^{\frac{-t^2}{2}}$$	α = intercept β = slope (≤18, ≥1 if restricted)	If the slope is allowed to be less than 1, the slope of the dose-response curve is infinite at zero dose. Φ is the standard normal density function, ϕ is the normal distribution function
Log-Probit $$p(d) = \gamma + (1-\gamma)\Phi[\alpha + \beta \log(d)]$$	γ = background (>0, <1) α = intercept β = slope (≤18, ≥1 if restricted)	For the log-probit model, the slope of the model will approach zero as dose approaches zero. However, depending on the data and parameter estimates, the slope for the log-probit model, for some relatively low doses perhaps less than those corresponding to the BMR, the slope can be quite steep, which may be manifested in terms of a relatively low value for the BMDL (or perhaps an "NA" result for the BMDL if this causes convergence problems because the steepness entails BMDL estimates that get very small).

Table 5. *Calculation of the BMD and BMDL for the individual dichotomous models*

Model	BMD Calculation	BMDL Calculation
Multistage	There is no general analytic form for the BMD in terms of the BMR and the estimated model parameters for the multistage model. Instead, the BMD is the root of the equation $\beta_1 BMD + \cdots +_n BMD^n + \ln(1 - A) = 0$, where $A = \begin{cases} BMR & extra\ risk \\ \dfrac{BMR}{1 - \gamma} & added\ risk \end{cases}$	The BMR equation is used as a non-linear constraint, and the minimum value of BMD is determined such that the log-likelihood is equal to the log-likelihood at the maximum likelihood estimates less $\dfrac{\chi^2_{1,1-2\alpha}}{2}$
Weibull	$BMD = \begin{cases} \left[\dfrac{-\ln(1 - BMR)}{\beta} \right]^{\frac{1}{\alpha}} & extra\ risk \\ \left[\dfrac{-\ln(1 - \frac{BMR}{1 - \gamma})}{\beta} \right]^{\frac{1}{\alpha}} & added\ risk \end{cases}$	To calculate the BMDL, the defining equations for the BMD are solved for the slope parameter β, which is then replaced in the original model equations. This makes **BMD** appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.
Gamma	Let $G(x; \alpha) = \dfrac{1}{\Gamma(\alpha)} \int_0^x t^{\alpha-1} e^{-t} dt$ be the incomplete Gamma function and $G^{-1}(\cdot; \alpha)$ be its inverse function. Then $BMD = \begin{cases} \dfrac{G^{-1}(BMR; \alpha)}{\beta} & extra\ risk \\ \dfrac{G^{-1}\left(\frac{BMR}{1 - \gamma}; \alpha\right)}{\beta} & added\ risk \end{cases}$	To calculate the BMDL, the defining equations for the BMD are solved for the slope parameter β, which is then replaced in the original model equations. This makes *BMD* appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.
Logistic	The BMD estimate for the Logistic model is defined implicitly by the following equation; an iterative numerical method is used to determine the value of the BMD $BMD = \dfrac{\ln\left(\frac{1-Z}{1+Z \times e^{-\alpha}}\right)}{\beta}$ where $Z = \begin{cases} BMR & extra\ risk \\ BMR \times \dfrac{1+e^{-\alpha}}{e^{-\alpha}} & added\ risk \end{cases}$	To calculate the BMDL, the defining equations for the BMD are solved for the intercept parameter α, which is then replaced in the original model equations. This makes **BMD** appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.
Log-Logistic	$\ln(BMD) = \begin{cases} \dfrac{\ln\left(\frac{BMR}{1 - BMR}\right) - \alpha}{\beta} & extra\ risk \\ \dfrac{\ln\left(\frac{BMR}{1 - \gamma - BMR}\right) - \alpha}{\beta} & added\ risk \end{cases}$	To calculate the BMDL, the defining equations for the BMD are solved for the intercept parameter α, which is then replaced in the original model equations. This makes **BMD** appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.

Model	BMD Calculation	BMDL Calculation
Probit	The BMD estimate for the Logistic model is defined implicitly by the following equation; an iterative numerical method is used to determine the value of the BMD $$BMD = \begin{cases} \dfrac{\Phi^{-1}\big(BMR[1-\Phi(\alpha)]+\Phi(\alpha)\big)-\alpha}{\beta} \ extra\ risk \\[2ex] \dfrac{\Phi^{-1}\big(BMR+\Phi(\alpha)\big)-\alpha}{\beta} \ added\ risk \end{cases}$$	To calculate the BMDL, the defining equations for the BMD are solved for the intercept parameter α, which is then replaced in the original model equations. This makes **BMD** appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.
Log-Probit	$$\ln(BMD) = \begin{cases} \dfrac{\Phi^{-1}(BMR)-\alpha}{\beta} \ extra\ risk \\[2ex] \dfrac{\Phi^{-1}\dfrac{BMR}{1-\gamma}-\alpha}{\beta} \ added\ risk \end{cases}$$	To calculate the BMDL, the defining equations for the BMD are solved for the intercept parameter α, which is then replaced in the original model equations. This makes **BMD** appear in the model equations as a parameter. For more details, refer to Section 9.5.5, "BMDL Computation," on page 50.

9.6 Bayesian Dichotomous Model Descriptions

The Bayesian dichotomous models used in BMDS 3.0 are identical to the frequentist parametric models described above. The main difference is that the Bayesian models incorporate prior information into the analysis, and this information is used in the model fit.

For datasets with a large number of observations, there should be little or no quantitative difference between the Bayesian fits and the frequentist fits.

When there are fewer data points, the prior will affect the fit. Here, the impact is most noticeable on hockey-stick dose-response curves and dose-response data that suggest strong supralinear fits. In these cases, fits are "shrunk back" to smoother dose-response relationships whose rate of change is less steep.

The models and priors are described in the following table.

Table 6. *The individual models used and their respective parameter priors. Note that* $\text{logit}(\gamma) = \log\left(\frac{\gamma}{1-\gamma}\right).$

Model	Constraints	Priors	Notes
Quantal linear $p_1(d) = \gamma + (1-\gamma)(1 - \exp[-\beta d])$	$\beta > 0$ $0 \le \gamma \le 1$	$\log(\beta) \sim \text{Normal}(0,1)$ $\text{liogit}(\gamma) \sim \text{Normal}(0,2)$	
Multistage $p_2(d) = \gamma + (1-\gamma)\left(1 - \exp\left[-\sum_i \beta_i d^i\right]\right)$	$\beta_l > 0$ $0 \le \gamma \le 1$	$\log(\beta_1) \sim \text{Normal}(0,0.25)$ $\log(\beta_i) \sim \text{Normal}(0,1)$ $i \ge 2$ $\text{logit}(\gamma) \sim \text{Normal}(0,2)$	Note the prior over the β_1 parameter expresses the belief that the linear term should be positive if the quadratic term is positive in the two hit model of carcinogenesis. For model averaging i = 2.
Weibull $p_3(d) = \gamma + (1-\gamma)(1 - \exp[-\beta d^\alpha])$	$\beta > 0$ $\alpha > 0$ $0 \le \gamma \le 1$	$\log(\beta) \sim \text{Normal}(0,1)$ $\log(\alpha) \sim \text{Normal}(\log(2),0.18)$ $\text{logit}(\gamma) \sim \text{Normal}(0,2).$	Here the prior over α is designed such that there is only a 0.01 prior probability the power parameter will be less than 1. This allows for models that are supra-linear; however, it requires a large amount of data for the α parameter to go much below 1.
Gamma $p_4(d) = \gamma + \frac{1-\gamma}{\Gamma(\alpha)} \int_0^{\beta d} t^{\alpha-1} \exp(-t)\, dt$	$\beta > 0$ $\alpha > 0.2$ $0 \le \gamma \le 1$	$\log(\beta) \sim \text{Normal}(0, 1)$ $\log(\alpha) \sim \text{Normal}(\log(2),0.18)$ $\text{logit}(\gamma) \sim \text{Normal}(0,2)$	Here the prior over α is designed such that there is only a low prior probability the power parameter will be less than 1. This allows for models that are supra linear; however, it requires a large amount of data for the parameter to go much below 1. The α parameter is also constrained to be greater than 0.2 for numerical reasons.
Dichotomous Hill $p_5(d) = \gamma + \frac{v(1-\gamma)}{1+\exp[-a-b\log(d)]}$	$0 \le \gamma \le 1$ $0 \le v \le 1$ $-\infty < a < \infty$ $b > 0$	$a \sim \text{Normal}(3, 3.3)$ $\log(b) \sim \text{Normal}(\log(2),0.5)$ $\text{logit}(\gamma) \sim \text{Normal}(-1,2)$ $\text{logit}(v) \sim \text{Normal}(0,3)$	
Logistic $p_6(d) = \frac{1}{1+\exp[-\beta_0-\beta_1 d]}$	$-\infty < \beta_0 < \infty$ $\beta_1 > 0$	$\beta_0 \sim \text{Normal}(0, 2)$ $\log(\beta_1) \sim \text{Normal}(0,1)$	
Log-Logistic $p_7(d) = \gamma + \frac{1-\gamma}{1+\exp[-\beta_0-\beta_1\log(d)]}$	$-\infty < \beta_0 < \infty$ $\beta_1 > 0$	$\beta_0 \sim \text{Normal}(0, 1)$ $\log(\beta_1) \sim \text{Normal}(\log(2),0.25)$ $\text{logit}(\gamma) \sim \text{Normal}(0,2).$	
Probit $p_8(d) = \Phi(\beta_0 + \beta_1 d)$	$-\infty < \beta_0 < \infty$ $\beta_1 > 0$	$\beta_0 \sim \text{Normal}(0,2)$ $\log(\beta_1) \sim \text{Normal}(0,1)$	
Log-Probit $p_9(d) = \gamma + (1-\gamma)\Phi[\beta_0 + \beta_1 \log(d)]$	$-\infty < \beta_0 < \infty$ $\beta_1 > 0$	$\beta_0 \sim \text{Normal}(0, 1)$ $\log(\beta_1) \sim \text{Normal}(\log(2),0.25)$ $\text{logit}(\gamma) \sim \text{Normal}(0,2)$	

9.6.1 Bayesian Model Averaging

From a Bayesian perspective, inference proceeds by defining a data generating mechanism given model M and its parameters. That is, one defines the likelihood $\ell(D|M)$ over the data D, which contains a data generating mechanism (e.g. normal, binomial), a model M, of the moments of that mechanism, and a prior probability model $\pi(M)$ over M. This probability model is $\pi(M) = f(M)f(\theta|M)$ where the prior probability of the model is $f(M)$ and $f(\theta)$ is a density over the model's parameters. For single model inference, $f(M) = 1$ and inference on model M given D is defined by Bayes theorem, which is

$$g(M|D) \propto \ell(D|M)\pi(M).$$

From a Bayesian perspective, functions of θ also have posterior densities, which are transforms of the original parameters. For example, the parameter of interest in this manuscript is the benchmark dose, which is a function of the original parameters, i.e., $h(\theta) = BMD$. In what follows, we define inference over K potential models and the BMD using Bayesian model averaging; here $0 \leq f(M) \leq 1$ and $\sum_{k=1}^{K} f_k(M_k) = 1$.

For a dataset D and a model M_k, we estimate each model M_k's parameter vector θ_k using maximum a posteriori estimation (MAP) and compute the posterior density of the BMD, i.e., $g_k(BMD|M_k, D)$. The posterior density of the model averaged BMD is

$$g_{ma}(BMD|D) = \sum_{k=1}^{9} \pi_k(M_k|D)g_k(BMD|M_k, D), \tag{1}$$

where π_k is the posterior probability of model M_k given the data. The BMD and BMDL are computed from this posterior density. More specifically, the point estimate of the BMD is the weighted average of individual MAP estimates and the BMDL estimate is taken as the $100(1 - \alpha)$ percentile for appropriately chosen probability level α. Model weights π_k are approximated using a Laplace approximation.

The posterior density of the individual models and the posterior model weights described in (1) are approximated using a Laplacian approximation. As described in the following sections, this approximation is similar to the Model Averaged Profile Likelihood (MAPL) approach of Fletcher and Turek (Fletcher and Turek, 2012). However, while MAPL relies only on the likelihood, our approach incorporates prior information in calculating the marginal profile density of the BMD. In other words, both the likelihood and prior are used. The model-specific density is defined by treating profile density bounds as quantiles of a marginal posterior density for the parameter of interest, and the relation to the present approach and the MAPL approach is justified asymptotically.

Our approach can be related to the MAPL framework by substituting the posterior density for the likelihood in each of the steps. Instead of information criteria-based weight, we use the Laplace approximation for model weights. This method approximates the marginal likelihood using the posterior MAP estimate and Hessian of the log-posterior.

For the model average approach, all Bayesian models described above are available in the model average. This is with the exception of the multistage model, which is capped to a maximum degree of 2. The reasoning for this follows upon the work of Nitcheva, et al. (2007) who show higher-order polynomials are not necessary and the fact that other models of the model averaging suite (e.g., dichotomous Hill) provide increased curvature.

9.6.2 Weight Calculation

In previous approaches to benchmark dose inference using model averaging, weights were calculated using either the BIC or AIC, where the AIC is used primarily in frequentist model averaging.

The proposed approach generates weights using the Laplace approximation to the marginal density of the data. That is, for model M_k, $1 \le k \le 9$, with parameter vector θ_k of length s, one approximates the marginal density as

$$I_k = (2\pi)^{S/2}|\hat{\Sigma}_k|^{1/2}\ell(D|M_k,\hat{\theta}_k)g(\hat{\theta}_k|M_k,D)\pi_k(M_k|D), \tag{2}$$

where

$\hat{\Sigma}_k$ is the negative inverse Hessian matrix evaluated at $\hat{\theta}_k$,

$\hat{\theta}_k$ is the MAP estimate,

$\ell(D|,M_k,\hat{\theta}_k)$ is the likelihood of the model given the data *D*, and

$g(M|D) = g(\hat{\theta}_k|M_k,D)\pi_k(M_k|D)$ is the MAP of the posterior density for M_k.

To compute the posterior model probabilities for M_k, one calculates the MAP and then calculates I_k using equation (2). The posterior probability of the model is

$$\pi_k(M_k|D) = \frac{f(M_k)I_k}{\sum_{i=1}^{9}f(M_k)I_k},$$

where $f(M_k)$ is the prior probability of model M_k (e.g., 1/9 if each of 9 models is treated as equally plausible *a priori*).

9.6.3 Computation of the Model-Averaged BMDL and BMD Point Estimate

Our model-averaged BMD point estimate is the weighted average of BMD MAP estimates from individual models, weighted by posterior weights $\pi_k(M_k|D)$. This is equivalent to the median of the approximate posterior density of θ. For the BMDL or BMDU estimates, equation (1) is integrated. A 100(*α*)% BMDU estimate or 100(1 - *α*)% BMDL estimate is the value BMD$_\alpha$ such that:

$$\alpha = \int_{-\infty}^{BMD_\alpha} g_{ma}(BMD|D) \, dBMD,$$

which is

$$= \sum_{k=1}^{9} \pi_k(M_k|D) \int_{-\infty}^{BMD_\alpha} g_k(BMD|M_k,D) \, dBMD. \tag{3}$$

In (3), the quantity $\int_{-\infty}^{BMD_\alpha} g_k(BMD|M_k,D) \, dBMD$ is approximated. That is,

$$\int_{-\infty}^{BMD_\gamma} g_k(BMD|M_k,D) \, dBMD$$

$$\approx \frac{1}{2}\Pr\left(-2\log[\hat{g}_k(\widehat{BMD}|M_k,D)] - 2\log[\hat{g}_k(BMD_\gamma|M_k,D)] < \chi^2_{1,\gamma}\right), \tag{4}$$

where $\hat{g}_k(x|M_k,D)$ is the maximum value of the posterior evaluated at *x,* where *x* is the MAP estimate BMD or the upper limit of integration BMD_γ, and $\chi^2_{1,\gamma}$ is the γ quantile of a chi-squared random variable with one degree of freedom. The above approximation assumes $BMD_\gamma < \widehat{BMD}$. When $\widehat{BMD} < BMD_\gamma$ the right-hand side of (4) is replaced by

$$\approx 1 - \frac{1}{2}\Pr\left(-2\,log[\hat{g}_k(\widehat{BMD}|M_k,D)] - 2\log[\hat{g}_k(BMD_\gamma|M_k,D)] < \chi^2_{1,\gamma}\right).$$

This approximation is like the profile likelihood used when estimating the BMDL and BMDU using the method of maximum likelihood, but in this case $\hat{g}_k(x|M_k,D)$ is the posterior density, which contains both the likelihood and the prior.

9.7   Frequentist Nested Model Descriptions

9.7.1 Special Considerations for Models for Nested Dichotomous Endpoints

The most common application of the models in this section will be to developmental toxicology studies of organisms that have multiple offspring per litter, as do rodents. In these study designs, pregnant females ("dams") are given one or several doses of a toxicant, and the fetuses, embryos, or term offspring ("pups") are examined for signs of abnormal development. In such studies, it is usual for the responses of pups in the same litter to be more similar to each other than to the responses of pups in different litters ("intra-litter correlation", or "litter-effect"). Another way to describe the same phenomenon is that the variance among the proportion of pups affected in litters is greater than would be expected if the pups were responding completely independently of each other.

The models in this section make available two approaches to this feature of developmental toxicology studies: they use a probability model that provides for extra inter-litter variance of the proportion of pups affected (the beta-binomial probability model: see the "Likelihood Function" section below); and they incorporate a litter-specific covariate that is expected to account for at least some of the extra inter-litter variance. This latter approach was introduced by Rai and Van Ryzin ([1985](#)), who reasoned that a covariate that took into account the condition of the dam before dosing might explain much of the observed litter effect. Those authors suggested that litter size would be an appropriate covariate. For the reasoning to apply strictly, the measure of litter size should not be affected by treatment; thus, in a study in which dosing begins after implantation, the number of implantation sites would seem to be an appropriate measure. On the other hand, the number of live fetuses in the litter at term would not be an appropriate measure if there is any dose-related prenatal death or resorption (this has apparently been ignored in most of the literature).

Carr and Portier ([Carr and Porter, 1991](#)), in a simulation study, warn that in situations in which there is no effect of litter size, statistical models that incorporate a litter size parameter, as do the models in BMDS, will often erroneously indicate that there is a litter size effect. Thus, the user should use litter size parameters with caution. Unfortunately, there are currently no good diagnostics for determining whether a litter size effect exists.

9.7.2 Likelihood Function

Let g represent the number of dose groups. For the ith group, there are n_i pregnant females administered dose $dose_i$. In the jth litter of the ith dose group there are s_{ij} **fetuses**, x_{ij} affected fetuses, and, potentially, a litter-specific covariate r_{ij} which will often be a measure of potential litter size, such as number of implantation sites, though this is not a requirement of the models. In what follows, the dose-response model, which gives the probability that a fetus in the jth litter of the ith dose group will be affected is represented by

$$p(dose_i, r_{ij})$$

The beta-binomial distribution can be thought of as resulting from sampling in two stages. First, each litter is assigned a probability, P_{ij} from a beta distribution (beta distributions represent a two-parameter family of probability distributions defined on the interval (0,1)). The parameters of the beta distribution are determined by the administered dose, the litter specific covariate r_{ij} and the degree of intra-litter correlation, v_i. Note that the intra-

litter correlation parameter varies among doses. It is well known (Williams et al., 1988) that when the true intra-litter correlation differs among doses, unbiased estimates of the other parameters in a dose-response model can only be obtained if dose-specific intra-litter correlation parameters are estimated. As a special case, if $v_i=0$, then this part of the process is completely deterministic, and

$$P_{ij} = p(dose_i, r_{ij})$$

This allows for the possibility of no litter effect at all.

In the second stage of sampling, *s_{ij} fetuses* are assigned to the litter, and the number of affected fetuses, *x_{ij}* is sampled from a binomial distribution with parameters *P_{ij}* and *s_{ij}*.

The log-likelihood function that results from this process is:

$$L = \sum_{i=1}^{g}\left\{\sum_{j=1}^{n_i}\left[\sum_{k=1}^{x_{ij}}\ln(p(dose_i, r_{ij}) + (k-1)\Psi_i)\right.\right.$$
$$\left.\left. + \sum_{k=1}^{s_{ij}-x_{ij}}\ln(1 - p(dose_i, r_{ij}) + (k-1)\Psi_i) - \sum_{k=1}^{s_{ij}}\ln(1 + (k-1)\Psi_i)\right]\right\}$$

Where

$$\Psi_i = \frac{\phi_i}{(1-\phi_i)}$$

And

$$\sum_{a}^{b}(\cdot) = 0$$

if *a > b* by convention.

9.7.3 Goodness of Fit Information—Litter Data

The "Litter Data" table provides a listing of the data, expected and observed responses and scaled residuals, for each litter.

The scaled residual values printed in the last column of the table are defined as follows:

$$(Obs. - Expected)/SE$$

where "Expected" is the predicted number of responders from the model and SE equals the estimated standard error of that predicted number. For these models, the estimated standard error is equal to sqrt[n*p*(1-p)*(*ϕ(n-1)+1)], where

- n is the sample (litter) size,
- p is the model-predicted probability of response, and
- f is the model-predicted intra-litter correlation coefficient.

The overall model should be called into question if the scaled residual values for several individual dose and litter-specific covariate combinations, particularly for the control group or a dose group near the BMD and for litter-specific covariate values close to the overall mean, are greater than 2 or less than -2.

The goodness-of-fit p-values are calculated using a bootstrap approach.

3. The MLE parameter values are used to generate B pseudo-datasets having the same design features (number of doses and number of litters per dose), litter-sizes, and, if necessary, litter-specific covariate values, as the original dataset. What varies from pseudo-dataset to pseudo-dataset are the number of responding "units" within litters, and those are generated, at random, as dictated by the values of the ML estimates.

4. Once the B bootstrap iterations are generated, a statistic referred to as *chi-square* is calculated for each. The chi-square statistic is the sum of the squares of the scaled residuals for each litter, as described above. Higher values of that statistic are indicative of poorer match between the model predictions and the data.

Note The chi-square statistic is so called here because, in traditional testing situations, that statistic would be approximated by a chi-squared random variable having a certain degree of freedom, and its "significance" (p-value) would be determined from the appropriate chi-squared distribution function.

5. The chi-square statistic from the original data is computed and compared to the values from the B bootstrap iterations. The p-value is the proportion of chi-square values from the iterations that are greater than the original chi-square value.

   High p-values are indicative of adequate fit (i.e., there was a high proportion of chi-square values associated with pseudo-datasets obtained from data known to be consistent with the model and the ML estimates of the model parameters).

That calculation is repeated three times, and various percentiles of the generated chi-square statistic are presented. This allows the user to determine if enough bootstrap iterations (B) have been specified. The default for B is 1000 and should probably not be reduced. The user may wish to increase the default if the percentiles for chi-square differ markedly across the three runs (in particular, the median and lower percentiles), or if the p-values calculated from the three runs differ markedly. This may only be an issue when the p-value is close to the value (e.g., 0.05 or 0.10) used as a critical value for deciding if the fit of the model to the data is adequate. If there is some variability in the p-values, but they are all greater than 0.20, for example, then one probably need not worry about increasing the value for B.

BMD Computation

BMD computation is similar to that for dichotomous models with the added wrinkle that a value for a litter-specific covariate (LSC) may be used, in addition to dose, to describe changes in the endpoint. It therefore affects the BMD calculation. If an LSC is included in the model, the user can choose to plot results and compute BMDs for one of two specific values of the LSC, either the overall mean (across all dose groups) or the control group mean. Typically, the overall mean is the preferred choice, but the control group mean might be appropriate in certain situations.

For example, suppose the LSC value varies enough from group to group to be "interesting," but it goes up for some dose groups and down for others in a manner that contraindicates a dose effect. In this case, you might decide to use the control group mean LSC when the BMD is close to the background dose (i.e., basically deciding that the LSC of interest in that region is more likely to be the average observed for the control group as opposed to the average across all the groups). If a covariate is found to be affected by dose, i.e., if its value appears to have a consistent trend with respect to dose, its use is discouraged.

BMDL Computation

BMDS currently only calculates one-sided confidence intervals, in accordance with current BMD practice. The general approach to computing the confidence limit for the BMD (called the BMDL here) is the same for all the models in BMDS, and is based on the asymptotic distribution of the likelihood ratio (Crump and Howe, 1985) .

The approach used for all the nested dichotomous models is the same. The equations that define the benchmark response in terms of the benchmark dose and the dose-response model are solved for one of the model parameters, using either the control group mean or the overall mean of the litter-specific covariate. The resulting expression is substituted back into the model equations, with the effect of re-parameterizing the model so that BMD appears explicitly as a parameter. A value for BMD is then found such that, when the remaining parameters are varied to maximize the likelihood, the resulting log-likelihood is less than that at the maximum likelihood estimates by exactly

$$\frac{\chi^2_{1,1-2\alpha}}{2}$$

9.7.4 Nested Dichotomous Frequentist Models

Table 6. *The individual dichotomous nested models used and their respective parameters*

Model	Parameters	Notes
Logistic Nested model $$p(d) = \alpha + \frac{\theta_1 r_{ij} + (1 - \alpha - \theta_1 r_{ij})}{(1 + e^{[-\beta - \theta_2 r_{ij} - \rho \times \ln(dose)]})}$$ if $dose > 0$, and $\alpha + \theta_1 r_{ij}$ if $dose = 0$	α = Intercept (≥0) ρ = power (≥0, can restrict ≥1) β = slope (≥0) θ_1 = first coefficient for the litter specific covariate θ_2 = first coefficient for the litter specific covariate $\phi_1, \dots, \phi_g$ = inter-litter correlation coefficients	in the model equation, r_{ij} is the litter-specific covariate for the j^{th} litter in the i^{th} dose group. In addition, there are g intra-litter correlation coefficients, $0 \leq \phi_i \leq 1$ ($i = 1, \dots, g$) $1 > \alpha + \rho \geq \theta_1 r_{ij} \geq 0$ for every r_{ij} If r_m represents either the control mean value for the litter-specific covariate or its overall mean, then the BMD is computed as: $$BMD = e^{\left\{ \frac{\left[\ln\left(\frac{A}{(1-A)} \right) - \beta - \theta_2 r_m \right]}{\rho} \right\}}$$ where $$A = \begin{cases} BMRF \ extra \ risk \\ \dfrac{BMRF}{(1 - a - \theta_1 r_m)} \ added \ risk \end{cases}$$ For the BMDL, the parameter β is replaced with an expression derived from the BMD definition and the BMDL is derived as described in Section 9.7.3.
NCTR model $$p(d) = 1 - e^{\{-(\alpha + \theta_1(r_{ij} - r_m) - (\beta = \theta_2(r_{ij} - r_m) \times dose^\rho\}}$$	α = Intercept (≥0) ρ = power (≥0, can restrict ≥1) β = slope (≥0) θ_1 = first coefficient for the litter specific covariate θ_2 = first coefficient for the litter specific covariate $\phi_1, \dots, \phi_g$ = inter-litter correlation coefficients	in the model equation, r_{ij} is the litter-specific covariate for the j^{th} litter in the i^{th} dose group, r_m is the overall mean for the litter-specific covariate $\theta_1(r_{ij} - r_m) \geq 0$ and $\theta_2(r_{ij} - r_m) \geq 0$ In addition, there are g intra-litter correlation coefficients, $0 \leq \phi_i \leq 1$ ($i = 1, \dots, g$) $1 > \alpha + \rho \geq \theta_1 r_{ij} \geq 0$ for every r_{ij} $$BMD = \left[\frac{-(\ln(1-A))}{(\beta + \theta_2 \delta_r)} \right] \times \left(\frac{1}{\rho} \right)$$ where $$A = \begin{cases} BMRF \ extra \ risk \\ \dfrac{BMRF}{(1 - a - \theta_1 \delta_r)} \ added \ risk \end{cases}$$ For the BMDL, the parameter β is replaced with an expression derived from the BMD definition and the BMDL is derived as described in Section 9.7.3.

9.8   Multi-tumor (MS_Combo) Model Description

The purpose of the MS_Combo program in BMDS is to allow the user to calculate BMDs and BMDLs for a combination of tumors (corresponding to a defined risk of getting one or more of those tumors) when the individual tumor dose-responses have been modeled using a Multistage-Cancer model.

Thus, the output of an MS_Combo run will present the results of fitting each individual tumor (including the BMD and BMDL for that tumor) plus the combined log-likelihood, BMD and BMDL for the combination of specified tumor responses.

In practice, the user should investigate each tumor individually and determine which degree of the Multistage-Cancer model is most appropriate for each individual tumor. That determination will involve all the usual considerations of fit, AIC, etc.

Once a specific form of the Multistage-Cancer model is chosen for each of the tumors of interest (they need not have the same degree across all the tumors in question), the user should specify those choices in the MS_Combo run.

Note   The following descriptions are valid only when the tumors are assumed to be independent of one another (conditional on dose level).

Because of the form of the multistage model, the MLE estimates for the combined risk are a function of the parameter values obtained for the individual tumor multistage model fits. In fact, the combined probability function has a multistage model form:

$$P(d) = 1 - e^{\{-(\beta_0 + \beta_1 d + \beta_2 d^2 + \cdots)\}}$$

and the terms of the combined probability function $(\beta_0, \beta_1, \ldots)$ are specified as follows

$$\beta_0 = \sum \beta_{0i}$$

$$\beta_1 = \sum \beta_{1i}$$

$$\beta_2 = \sum \beta_{2i}$$

etc.

where the sums are over i = 1, …, t, with

t being the number of tumors under consideration, and

β_{xj} being the x^{th} parameter (0, 1, …) for tumor j.

The β_{xj} values are available directly from the Multistage-Cancer runs performed on the individual tumors, but MS_Combo performs the calculations for the user, completing the summations of the individual terms and computing the BMD based on the combined parameter values and the user-specified BMR.

A profile-likelihood approach is used to derive the BMDL.

1. Given the BMD and the log-likelihood associated with the MLE solution, a target likelihood is defined based on the user-specified confidence level (e.g., 95%).
2. That target likelihood is derived by computing the percentile of a chi-square (1 degree of freedom) corresponding to the confidence level specified by the user (actually, the alpha associated with the confidence level, times 2).
3. That percentile is divided by 2 and subtracted from the maximum log-likelihood.

4. That derivation is based on a likelihood ratio test with one degree of freedom; it can be shown that estimating the BMDL corresponds to losing one degree of freedom, regardless of the number of tumors being combined.

The BMDL for the combined response (one or more of the tumors of interest) is defined as the smallest dose, D, for which the following two conditions are satisfied:

5. There is a set of parameters such that the combined log-likelihood using D and those parameters is greater than or equal to the target likelihood), and
6. For that set of parameters, the risk at D is equal to the user-specified BMR.

Note that the combined log-likelihood is a function of the fits of the individual tumors (the sum of the individual log-likelihoods), obtained using their tumor-specific β values. Thus, the search for the parameters of the combined Multistage-Cancer model varies the individual-tumor β values in such a way that the individual log-likelihoods add up to a combined likelihood within the range desired (greater than or equal to the target). However, to satisfy the second constraint, the sums of the individual-tumor parameters (shown above to be the parameters of the combined probability function) are used to evaluate the risk for any proposed BMDL, D.

Note that the individual tumors need not be modeled with the same degree of the Multistage-Cancer model. Any terms not included for an individual tumor are assumed to be zero (and will remain at zero during BMDL optimization) in the summations shown above. The optimizer DONLP2 is used for the combined BMDL estimation.

10.0 Troubleshooting

10.1   Avoid Using Windows Reserved Characters in File and Path Names

BMDS allows any character, except for Windows reserved characters, to be used when naming files or directories that BMDS will access.

However, the following Windows reserved characters are still disallowed and cannot be used:

< (less than)

> (greater than)

| (vertical bar or pipe)

? (question mark)

* (asterisk)

" (double quote)

Note   The backslash (\) should be used when specifying network drive paths.

10.2   Request Support with eTicket

For any technical problem related to running BMDS, please submit a problem report at the BMDS eTicket site. With eTicket, you can request help, ask a question, or check on the status of an existing issue.

11.0 References

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In BN Petrov; F Csaki (Eds.), 2nd International Symposium on Information Theory (pp. 267-281). Budapest, Hungary: Akadémiai Kiadó.

Carr, GJ; Porter, CJ. (1991). An evaluation of the Rai and Van Ryzin Dose-Response Model in teratology. Risk Anal 11: 111-120. http://dx.doi.org/10.1111/j.1539-6924.1991.tb00581.x

Crump, KS. (1984). A new method for determining allowable daily intakes. Toxicol Sci 4: 854-871. http://dx.doi.org/10.1016/0272-0590(84)90107-6

Crump, KS; Howe, RB. (1985). A review of methods for calculating statistical confidence limits in low dose extrapolation. In DB Clayson; D Krewski; I Munro (Eds.), Toxicological risk assessment: Vol 1 Biological and statistical criteria (pp. 187-203). Boca Raton, FL: CRC Press, Inc.

Fleiss, JL; Levin, B; Paik, MC. (2003). Statistical methods for rates and proportions. Hoboken, NJ: John Wiley & Sons, Inc.

Fletcher, D; Turek, D. (2012). Model-averaged profile likelihood intervals. Journal of Agricultural, Biological, and Environmental Statistics 17: 38-51. http://dx.doi.org/10.1007/s13253-011-0064-8

Jeffreys, H. (1998). The theory of probability. Oxford, United Kingdom: Oxford University Press.

Kimmel, CA; Gaylor, DW. (1988). Issues in qualitative and quantitative risk analysis for developmental toxicology. Risk Anal 8: 15-20. http://dx.doi.org/10.1111/j.1539-6924.1988.tb01149.x

Nitcheva, DK; Piegorsch, WW; West, RW. (2007). On use of the multistage dose-response model for assessing laboratory animal carcinogenicity. Regul Toxicol Pharmacol 48: 135-147. http://dx.doi.org/10.1016/j.yrtph.2007.03.002

Rai, K; Van Ryzin, J. (1985). A dose-response model for teratology experiments involving quantal responses. Biometrics 41: 1-9.

RIVM (National Institute for Public Health and the Environment (Netherlands)). (2018). PROAST. Retrieved from https://www.rivm.nl/en/Documents_and_publications/Scientific/Models/PROAST

U.S. EPA (U.S. Environmental Protection Agency). (2012). Benchmark dose technical guidance. (EPA/100/R-12/001). Washington, DC: U.S. Environmental Protection Agency, Risk Assessment Forum. https://www.epa.gov/risk/benchmark-dose-technical-guidance

Appendix A: Version History

The following sections document new, changed, or updated BMDS features as documented in each version's respective readme file. The information is included here for historical and reference purposes.

A.1   BMDS 1.2

September 2, 2000 - A new user interface (BMDS0900.exe) was distributed to fix some problems with installation of BMDS on certain Windows 98 configurations. If you successfully installed BMDS version 1.2 using a previous installation procedure you do not need this upgrade. This upgrade merely simplifies the installation process and corrects some problems that did not allow BMDS to install to certain computer hardware/software configurations. (This version of the software is no longer being made available as there are newer versions now available which fix problems that were being encountered on newer operating systems.)

A.2   BMDS 1.2.1

October 25, 2000 - A new version of BMDS, version 1.2.1, is being distributed at this time. This version contains new versions of the continuous Polynomial (version 2.1) and Hill (version 2.1) models. If you do not want to completely reinstall BMDS, you can download the new model executables and run them separately or under the BMDS version 1.2 interface. These new versions of the polynomial and Hill models fix problems associated with running the model on Windows NT/2000 operating systems, provide improved model fit for certain unique datasets and improve upon the rate of convergence on a BMD and BMDL.

A.3   BMDS 1.3

March 22, 2001 - Version 1.3 of BMDS is now available! This latest version of BMDS, version 1.3, contains new continuous Polynomial (v2.1), Power (v2.1) and Hill (v2.1) models, new dichotomous Multistage (v2.1), Weibull (v2.1) and Gamma (v2.2) models, and an improved user interface. The new models are more compact and stable (will converge on BMD and BMDL solutions more often). The user interface upgrades are described in the new help manual (PDF format) for version 1.3 and the readme.txt file that is distributed with the upgrade.

A.4   BMDS 1.3.1

January 22, 2002 - Version 1.3.1 of BMDS is now available! Version 1.3.1 contains a revised help manual and user interface, including a revision to the interface that allows the Multistage model to calculate BMD and BMDL values for very low (below E-5) benchmark response (BMR) levels.

November 13, 2002 - A new polynomial model (Version 2.2) is now available that fixes the previous incompatibility with Windows 2000. Download it to your main bmds directory (same directory as the bmds.exe file).

A.5   BMDS 1.3.2

May 23, 2003 - Version 1.3.2 of BMDS is now available! Version 1.3.2 contains revised polynomial (poly.exe) and nested logistic (nlogist.exe) models that are compatible with Windows 2000. If you are using a Windows 98 or older operating system, you may need to update your msvcrt.dll driver. We suggest that you obtain the latest msvcrt.dll driver from Microsoft or download this version of the msvcrt.dll driver and copy it to the c:\windows\system directory of your computer (you may have to exit Windows and do this in DOS mode).

A.6   BMDS 1.4.1

February 5, 2007 - Version 1.4.1 is now available! All models have been recompiled to improve speed, stability and compatibility with the latest Windows operating systems. Improvements have been made to the model output format for all models. A Multistage-Cancer model has been added that calculates and reports a cancer slope factor and plots the linear extrapolation from the BMDL to the background response estimate per EPA's 2005 cancer guidelines. Unlike the Multistage model, the Multistage Cancer model does not estimate added risk, nor does it allow beta coefficients to be unrestricted. The Quantal-Quadratic model was removed from Dichotomous model choices (note: the user can still run this model by specifying the power term of the Weibull model to be 2, but this model is not retained in the BMDS dichotomous model listings)

Issues in the continuous models that caused occasional errors in degrees of freedom assignments which impacted continuous model test results have been resolved. Acceptance criteria for Tests 2, 3 and 4 was changed from $p>=0.05$ to $p>=0.1$ and default risk type changed to "Std. Dev." for all continuous models to be consistent with EPA's draft BMD technical guidance (EPA, 2000). Issues with the Hill model have been fixed, including memory problems which were causing some operating systems to crash. Parameter standard error estimates and Chi-squared residual calculations in all the continuous models were checked and corrected if in error. Model A3 of the continuous model testing procedures has been modified so that it always uses the user-specified value for the parameter rho, including the constant-variance case where rho = 0. When rho = 0, model A3 is the same as model A1, and it is reported explicitly in the constant-variance runs. As a consequence, all model runs report the entire set of models (A1, A2, A3, R and the fitted model) and all four hypothesis tests.

Issues in the Nested models that caused occasional errors in degrees of freedom assignments have been resolved. Memory problems which were causing problems for some NCTR model runs have been fixed.

August 29, 2007 - BMDS Version 1.4.1b has been added to replace version 1.4.1. This version contains an update to the BMDS help file.

November 9, 2007 - BMDS version 1.4.1c is now available. This version updates dichotomous models that were already included on BMDS version 1.4.1b. The updates primarily improve the handling of parameter specifications, particularly in situations where the user may wish to specify the background parameter to be zero.

A.7   BMDS 2.0 (beta)

September 28, 2007 - BMDS Version 2.0 beta is now available for inspection and testing (NOTE: this is a beta test version, provided only for your examination and testing - BMDS 1.4.1b should be used for definitive risk assessment calculations). BMDS 2.0 beta employs a new graphical user interface and makes it easy to run a number of models for

one dataset and compare the results. BMDS 2.0 beta also has a new set of quantal models with alternative background parameters (i.e., background additive to dose). We welcome comments and suggestions on the functioning of the interface and its new features, and on the new models.

October 10, 2007 - BMDS 2.0 beta - Build 19 released on October 10, 2007 replaces the first BMDS 2.0 beta release of September 28, 2007 (Build 13). The new Build 19 has important changes and enhancements as a result of additional testing and user exposure and should be downloaded and used instead of Build 13. Enhancements include the ability to better run a number of the BMD models and also added flexibility and fixes for user interface features. Changes include the designation of the new Dichotomous models as Alternate Dichotomous to better reflect their production status. Please refer to the readme.txt file included with the software installation for more details on the BMDS 2.0 beta.

A.8 BMDS 2.0 (final)

July 10, 2008 - BMDS Version 2.0 final is now available. Released on July 10, 2008, it replaces BMDS 1.4.1c as the official BMDS software. BMDS 2.0 is a rewrite of the user interface and risk assessment modeling framework, with a markedly improved functionality and enhanced multi-model processing capabilities. It uses the same underlying source code for the models in BMDS 1.4.1 software, with minor corrections and some important additions. For details on the new user interface, go to the BMDS 2.0 Help menu option in the installed software. BMDS 2.0 also has a new set of quantal models with alternative background (i.e., background additive to dose) and asymptote (i.e., Hill model) parameters, as well as a Beta Exponential set of models.

A.9 BMDS 2.1 (beta)

September 30, 2008 - EPA is making version 2.1 of BMDS available at this time for public beta testing. Version 2.1 includes a beta (external peer review) version of a new time-dependent toxicodiffusion model for continuous outcomes (Zhu et al., 2005), incorporates graphical plots for the continuous exponential models and allows for the use of individual animal continuous response data. The BMDS toxicodiffusion model was developed by the USEPA National Center for Environmental Assessment (NCEA), through partnerships with the USEPA Neurotoxicology Division (NTD) and the University of South Florida, to characterize toxic effects (e.g., neurotoxicity) that potentially evolve along critical time points. It does this by:

- modeling a dose-response along a time-course of repeated response measures; and
- computing benchmark doses and their confidence limits along the time course.

Documentation for the toxicodiffusion model can also be downloaded. The documentation contains a full description of the model, input requirements, model run options and sample runs.

In addition, EPA is distributing an external review (beta) version of a concentration-time (CxT) model originally programmed by Wil ten Berge. The EPA ten Berge model implements an approach to evaluating the CxT relationships for effects associated with chemical exposures. The EPA's version 1.0 implementation of this model is being distributed along with associated documentation and comments on the model received from external peer reviewers. EPA plans to respond to external review comments and incorporate the ten Berge model into a future version of BMDS.

A.10   BMDS 2.1 (Build 52)

July 30, 2009 - EPA is now distributing the final release of Version 2.1 (Build 52) of the Benchmark Dose Software (BMDS). BMDS 2.1 (Build 52) contains user interface enhancements as well as several additions/enhancements to the suite of models available for modeling dose-response data, including new features for the continuous exponential models and a new interface for the ten Berge concentration-time model. For details on the changes to the user interface, go to the BMDS 2.1 Help menu option in the installed software. The Readme.rtf file distributed with BMDS describes the improvements made in version 2.1 (Build 52), installation requirements, and known problems.

The exponential models contained in this version of BMDS have been developed in conjunction with the Netherlands' National Institute for Public Health and the Environment (RIVM) to be consistent with the exponential models contained in the RIVM's PROAST software . The US-EPA and RIVM are working together to achieve consistency between the BMDS and PROAST software and methods.

A.11   BMDS 2.1.1 (Build 55)

November 9, 2009 - EPA is now distributing Version 2.1.1 (Build 55) of the Benchmark Dose Software (BMDS). BMDS 2.1.1 (Build 55) contains a flexible new feature that allows users to export select BMDS summary report data and plots to Excel. It also contains a comprehensive set of sample session and model option files to assist users in running batch operations, and several improvements to the ten Berge model that were not available in version 2.1. The Readme.rtf file distributed with BMDS provides details on the improvements made in Version 2.1.1 (Build 55), installation requirements, and known problems.

A.12   BMDS 2.1.2 (Build 60)

June 11, 2010 - BMDS 2.1.2 (Build 60) contains user interface enhancements to the "Summary Report" feature, new sample session and model option files, and improvements to the ten Berge model that were not available in version 2.1.1:

- BMDS 2.1.2 can access folders or files with embedded space(s) in them. Combinations of the path and file name must be less than 256 characters.
- The default location for searching for certain files (i.e. "Session," "Option," "Data," "Plot," and "Output" files) is now the last location (folder) to which the user saved that type of file or from which he accessed that type of file. There is a separate "memory" for file location for each of the above-listed file types.
- Improvements have been made to the "Export to Excel" feature associated with the "Summary Report" of a session run.. This feature allows the user to select which variables will be exported to Excel by checking/un-checking, in the "Export to Excel" column, the boxes corresponding to the variable rows the user wishes to export. The plots are exported as well, in a separate Excel "Plots" sheet.
- Additional session and option files have been added to the "SessionFiles" and "OptionFiles" folders. These folders contain sample sessions that allow the user to quickly run a specific set of models and model options for a selected dataset.
- The dichotomous Weibull model can accept a lower bound on power specified by the user. The EPA default choice (power restricted to be greater than or equal to 1) can still be checked, but if the user wishes not to restrict the power in that way, s/he may specify a value greater than or equal to zero; zero was the only other option in

previous versions and that could cause biologically unrealistic fitted curves and numerical problems.

- The Toxicodiffusion model no longer needs the R(D)com Active-X control and will run with any version of R, 2.6.2 or later.

- Plots created by BMDS can be viewed using capabilities within BMDS or using GnuPlot. Those two options are accessed under the "Tools" menu item, the "View Plot" option. Using GnuPlot, the user can edit the plots to choose colors, fonts, line styles, etc.

- The option screen for the ten Berge CxT (concentration-time) model BMDS 2.1.2 had been improved to – address some issues and to allow saving the output file with a name of the user's choosing. More importantly, CxT results from the ten Berge model are automatically exported to an Excel template file containing two customizable plots, one that shows the contour of concentration and time combinations for a user-specified probability and a second that shows the probability of response as a function of either concentration or time. In both plots, the user specifies the values of the other variables in the model.

Like BMDS 2.1.1, BMDS 2.1.2 contains the following additions/enhancements over BMDS 2.0:

- The ten Berge CxT model alluded to in item 5 above, allows for fitting of dichotomous response datasets having two or more explanatory variables (as in acute inhalation toxicity experiments). The explanatory variables can be entered as main effects or in interaction (cross-product) terms. The user can request the value (and its bounds) of one explanatory variable when a response rate is specified (fixing the other explanatory variables at some user-specified values) and/or conversely, the value (and its bounds) of the response rate, given specification of all explanatory variables.

- The executable for the set of models known as the exponential models, proposed by Dr. Wout Slob of RIVM in The Netherlands has been expanded to allow the assumption of log-normally distributed data (the previous versions of the exponential models and all other continuous models in BMDS assume that the data are normally distributed). The four exponential models fit by BMDS are defined and labeled as follows:

  - Model 2: $m(dose) = a*\exp\{sign*b*dose\}$
  - Model 3: $m(dose) = a*\exp\{sign*(b*dose)d\}$
  - Model 4: $m(dose) = a*(c – (c-1)*\exp\{-1*b*dose\})$
  - Model 5: $m(dose) = a*(c – (c-1)*\exp\{-1*(b*dose)d\})$

  where "sign" indicates the direction of change in the responses (sign=+1 for increasing responses; sign=-1 for decreasing responses).

- A version of a new ToxicoDiffusion model for continuous outcomes (Zhu et al., 2005) allows for the analysis of repeated-measures data. The BMDS Toxicodiffusion model is able to characterize toxic effects (e.g., neurotoxicity) that potentially evolve with time points by

  - Modeling a dose-response along a time-course of repeated response measurements;
  - Computing benchmark doses and their confidence limits along the time course.

  The ToxicoDiffusion model includes graphical outputs showing the observed and model-predicted time-course data, residuals, and a summary of the bootstrap-based BMDL calculations.

For further detail see Zhu, Y., Jia, Z., Wang, W., Gift, J., Moser, V.C., and B.J. Pierre-Louis (2005), *Data Analysis of Neurobehavioral Screening Data: Benchmark Dose Estimation. Regulatory Toxicology and Pharmacology*, pp 190-201.

A.13   BMDS 2.2 (Build 66)

September 6, 2011 - EPA is now distributing the final release of Version 2.2 (Build 66) of the Benchmark Dose Software (BMDS). Key enhancements in BMDS 2.2 (Build 66) include:

- **Multiple Tumor Analysis -** BMDS 2.2 adds the capability to perform a combined analysis of multiple tumors. If the user is willing to assume that those tumors are independent and are well described by a multistage-cancer model, then the Multiple Tumor Analysis capability (accessed through the File/New or File/Open tool-bar choices) allows the user to estimate BMDs and BMDLs for the combined incidence of the tumors in question (i.e., BMDs and BMDLs for the likelihood of getting one or more of those tumors).
- **Trend Test for Dichotomous Data -** Another major addition is the new capability to perform a trend test on dichotomous datasets. This is the first in a series of trend test to be added to BMDS (future versions will also include trend test for continuous and nested data). The trend testing feature can be found on the dataset screens, accessible once a dataset has been identified by the user as containing dichotomous response data. The test performed is the Cochran-Armitage trend test described by Haseman (1984).
- **The Dichotomous Hill model has been modified -** Changes to the parameter initialization section of the Dichotomous Hill code have improved the convergence features of this model.
- **Automatic Transfer of Variable Name Changes to Other Option Files in a Session -** When working within a session, variable name changes (e.g., for dose, sample size, response, mean, or standard deviation variables) made in one option file (i.e., for one model) can be "transferred" to other option files included in that session (i.e., those for other models). The user will be prompted to determine if variable name assignment changes made in one option file should be made in all other option files included in that session. Thus, users can change variable name assignments once in a session, without having to make those changes separately in every option file.
- **Default Column Headers for New Datasets -** Note also that newly created dichotomous, continuous or nested model data files will start with default column headers, in a particular order, as appropriate for the type of data (e.g., Dose, N, and Effect for dichotomous datasets; Dose, N, Mean, and Std for summarized continuous datasets). The user may change those default headers, but will be warned that doing so may affect the running of BMDS-supplied sessions that look for those default names.

A.14   BMDS 2.2 (Build 67)

December 8, 2011 - EPA is now distributing BMDS 2.2 (Build 67), which include minor modifications to the user interface.

A.15   BMDS 2.3 (Build 68)

September 12, 2012 - BMDS 2.3 introduces more flexible error-trapping functionality, enabling users to store essential comments, documentation, notations, etc. on their data in the datasets and spreadsheets without triggering a data validation error.

Previously, BMDS ran validation checks on the dataset each time the file was opened. BMDS would interpret notes, comments, extraneous text, etc. contained in the dataset as

data errors and would not open the file. This would require opening the file in a separate program and purging any comments or text.

In BMDS 2.3, the application instead checks for data errors when the data is sent to the model. BMDS bases its validation on the dataset's column assignments from the Option Screen.

After you open the dataset, select the Model and Model Type, and then select Proceed, the Option Screen displays. In the Column Assignments section of the Option Screen, specify the data labels corresponding to the dataset variables**.**

When you click Run, BMDS will validate and error-check the data for only those assigned columns. BMDS will then display a message box describing any errors related to invalid values or blank cells found in those columns.

BMDS will also flag any duplicate column headers if they conflict with the Column Assignments specifications. Note that BMDS is case-sensitive, so BMDS considers dose, Dose, and DOSE to be separate variable names.

BMDS 2.3 also features the following fixes and enhancements:

- Fixes a problem reported by users of Microsoft Office in Windows 7, in which clipboard errors would crash both BMDS and Microsoft Office.
- Warns the user when saving a file using a name that exceeds the Windows limit of 256 characters. To correct the problem, rename the file or move the file to a directory higher in the hierarchy.
- Enhances Option screen validation to provide more thorough parameter constraint checking and to display pop-up messages describing errors.
- Fixes bug that prevented BMDS from adding a new/existing dataset to the Session Grid.
- Fixes bug related to dynamically drawn dialog boxes.
- An issue was discovered in the computation of the A3 model log-likelihood for the continuous models when the user specified the variance parameter alpha. While this issue is being investigated and resolved, the option to specify alpha has been disabled for those models.
- Fixes intermittent bug where occasionally, when running the exponential model, BMDS failed to display a summary graph and report, even when output files were produced.
- Edits and updates to several Help topics:
  - Updated all Model Description topics to reflect upper parameter limit of 18 for some models.
  - Added the following topics under "Other Data or Analysis Types": Data With Negative Means, Test for Combining Two Datasets for the Same Endpoint
  - Added the following Troubleshooting topics: Help File Does Not Display, Decimal Separator Should Be a Period (see Section 3, "Usage Tips," in this file for more information)
  - Added description of multi-tumor model (.d) file to the Multiple Tumor Analysis topic.
  - Added text to the Dichotomous Hill and Logistic models descriptions to address occasional error messages that appear for these models.
  - Under the "Graphic Output from Models" topic, added the subtopic "Error Bar Calculations," explaining how BMDS generates error bars for various model types.

A.16   BMDS 2.3.1 (Build 69)

September 28, 2012 - BMDS maintenance release Version 2.3.1 (Build 69) replaces Version 2.3 (Build 68). It simplifies how data validation errors for certain models are reported to the user. BMDS 2.3 also introduced user interface improvements.

None of the actual dose-response models in BMDS were modified for versions 2.3 and 2.3.1.

A.17   BMDS 2.4 (Build 70)

BMDS 2.4 includes several enhancements to improve usability and ensure accurate and reliable results, such as more informative plot titles and the ability to change an option file on the fly within a session. The help file includes a new section on "BMDS Best Practices" containing valuable information and guidance on such topics as optimization criteria, alternative models, re-initializing parameters, and lognormal response option, among others.

Also, in this release, the BMDS install package includes ICF International's BMDS Wizard, an Excel-based tool that facilitates the preparation and organization of, and enhances the reporting capabilities of, BMDS modeling sessions.

BMDS 2.4 adds the following new or enhanced functionality:

- BMDS plots now provide more informative titles, such as "Weibull Model, with BMR of 10% Extra Risk and 0.95 Lower Confidence Limit for the BMD (BMDL)."
- BMDS now includes the cancer slope factor for the MS_Combo output.
- When an option file is modified via the Session Grid and saved under a new name, BMDS now automatically links that file to the current session.
- In a modeling session, after changes are made to a model option file, BMDS lists the changes for the user and asks whether they should be applied to the other model option files in the session. A warning is given to remind the user that changes to option files will affect other sessions that use those option files.

BMDS 2.4 also features the following fixes or changes:

- Gnuplot files now remain on the screen until the user closes them.
- Fixes a problem where exporting rows to Excel would shift column values to different headers on the Dichotomous Format and Continuous Format worksheets.
- Fixes a problem where BMDS dichotomous models treated "%Positive" as Incidence data.
- Fixes a problem where cut and paste operations on session data fields worked erratically on Windows 7 systems.
- Fixes a problem where BMDS could not open datasets saved with a capital .DAX extension.
- Fixes a problem where selecting a column in the data grid for a Log10 transformation returned incorrect results.
- Removes "Exact" as an available solution in the exponential model for summary data.
- Removes the "Extra" option for all continuous models but Hill.
- Removes the Optimization section from the options screens for all models that are not affected by optimization settings.
- Changes the default model iterations for all models from 250 to 500.
- Harmonizes the output for all continuous models so that when the "Rel Dev." BMR type is selected in the option screen, the "Risk Type" reported in the output file says "Relative deviation" rather than "Relative risk."

- Fixes a problem in the Multi-tumor Model (ms_combo) where BMDS failed to return a valid combined BMDL if the largest dose in the first dataset listed in the tumor analysis was less than the maximum dose in the other tumor analysis datasets.
- Edits and updates several Help topics, including:
- A new section, "BMDS Best Practices for Obtaining Optimal Model Convergence."
- A new topic under Troubleshooting: "Avoid using special characters in filenames."
- A new topic documenting the model files (with their version numbers) used in BMDS 2.4.
- A new appendix for BMDS version history information that originally appeared in those versions' readme files.
- A new note on log transformations for the topic "Data Transformation Types."

ICF INTERNATIONAL'S BMDS WIZARD

The BMDS 2.4 install package includes ICF International's BMDS Wizard, an Excel-based tool that facilitates the preparation and organization of and enhances the reporting capabilities of BMDS modeling sessions. This install includes multiple copies of the ICF BMDS Wizard files that are preformatted for continuous, dichotomous, and dichotomous-cancer datasets.

BMDS "power users" employ ICF BMDS Wizard as a shell to simplify the BMD modeling process by streamlining data entry, model selection, option file development, output file reporting, and model comparisons.

ICF BMDS Wizard 1.7 can also export Microsoft Word-formatted reports that employ the latest EPA-approved reporting format (as of February 15, 2013). It can only export reports for continuous, dichotomous, and dichotomous-cancer models.

To run the ICF BMDS Wizard, go to the BMDS 2.4 program directory and locate the "BMDS Wizard 1.7" subdirectory. ICF has included a readme and quickstart guide to the software. Please refer to that documentation for details on running the tool.

More information on the ICF BMDS Wizard can be found at ICF's site: http://www.icfi.com/insights/products-and-tools/bmds-wizard.

Please note that ICF BMDS Wizard is not endorsed or approved by EPA. Please contact ICF International at wizard@icfi.com for support.

A.18   BMDS 2.5 (Build 82)

BMDS 2.5 provides updates to ICF International's BMDS Wizard, which includes support for 32-bit and 64-bit versions of Microsoft Office on Windows 7, as well as a new MS_Combo (multi-tumor) template.

In addition, BMDS 2.5 contains various improvements to model stability and reliability.

Resolved Issues

BMDS 2.5 features the following enhancement and fixes:

- The BMDS Tools>View Plot menu entry has been simplified to make it easier to generate a plot from a previously created .plt file.
- The Multistage, Multistage Cancer and MS_Combo models now provide accurate results when non-integer input data values for Incidence and Number of Subjects used.
- The Multistage, Multistage Cancer and MS_Combo models now work correctly when beta parameters are specified by the user. Previously, the models would fail to calculate BMDL.

- The Power model now honors the direction of adversity specified by the user. Previously, the model always determined the direction automatically.
- An intermittent crash in the MS_Combo model has been resolved.
- Several Help topics have been revised, including:
- Specified in the Troubleshooting section that path+filenames should not exceed 127 characters.
- Added a topic documenting the MS_Combo model input file.
- Added Appendix B, which documents, for each model options screen, all of that screen's fields, including any field constraints.
- Updated the following topics to reflect BMDS .out files content: Continuous Model Text Output, Continuous Model Maximum Likelihood, Tests of Fit.

A.19 BMDS 2.6

BMDS 2.6 includes several enhancements to improve usability and ensure accurate and reliable results.

BMDS 2.6 also introduces a significant new feature: the ability for BMDS to automatically detect and, optionally, install software updates. This feature will ensure BMDS users have access to the latest version of the software with up-to-date fixes and enhancements.

The BMDS install package includes ICF International's BMDS Wizard, an Excel-based tool that facilitates the preparation and organization of, and enhances the reporting capabilities of, BMDS modeling sessions. See the next section on the ICF BMDS Wizard for more information.

BMDS 2.6 features the following enhancement and fixes:

- Resolved the following issues related to the nested models (NLogistic, NCTR, and Rai and van Ryzin):
    - The approach to evaluating goodness-of-fit for the nested models has been changed. Now, a bootstrap-based approach (using the fitted model and the underlying beta-binomial distribution of the observations) is used to evaluate the lack of fit. This obviates the need to group the litter-specific observations across litters and avoids asymptotic approximations.
    - Fixed a calculation error that occurred when the litter-size covariate (LSC) was larger than the litter size. (Although the reliability of the Rai and van Ryzin model has been improved, there is a remaining issue where the model erroneously reports the same value for the BMD and BMDL under certain circumstances.)
    - Added scaled residual of interest calculations to the results report. Output includes min, average, and max scaled residual of interest, and number of litters used for the calculation.
- BMDS now handles file and path names more robustly across all models. Specific changes include:
    - Full path length (folder plus file name) is now 255 characters rather than 127.
    - Spaces and ampersands (&) are now permitted.
    - Fixed several issues that occurred when the user specified output and session names in some models.
    - BMDS now supports UNC (network) path names structured as *ComputerName**ShareName**Path* (e.g., *FileSrv1**Users**JDoe**USEPA**BMDS260*). The total number of characters cannot exceed 255.

- The ToxicoDiffusion and MS_Combo models are now out of "beta" status. Detailed technical documentation of these models can be viewed or downloaded from the BMDS Download page.
- The ToxicoDiffusion model can now run when R has been installed on a per-user basis.
- Resolved an issue with the Linear and Polynomial models where incorrect MLE's were produced for non-constant variance in some situations.
- Resolved an issue so that BMDS now correctly calculates and displays Standard Error Estimates for all models in the "Parameter Estimates" table, whereas previously most models displayed "*".
- Changed the color scheme for BMDS plots so they are easier to read.
- Included the most recent version of gnuplot (version 4.6.3).
- Fixed the Export to Excel function so it can export results for sessions containing any number of models. (Previously, sessions containing more than 24 models produced an error during export.)
- Fixed an issue so that BMDS now supports regional settings for the decimal separator in the user interface and in spreadsheets created by the Export to Excel function. However:
  - Files generated by BMDS (such as the .out files and .d files) will still show "." as the decimal.
  - No "thousands separator" (regardless of regional setting) can be used in the data; that is, one thousand can only be written as 1000 rather than as 1,000.
- Added the following items to the summary table in the BMDS user interface: "Scaled Residual for Control Group" and "Scaled Residual for Dose Group Nearest the BMD". The relevant values are already printed out in the "Goodness of Fit" table.
- Improved the e-Ticket system for users to request support and guidance.
- Fixed an issue in the Power model where the specified power parameter output value was displayed incorrectly in the output text.
- Fixed an error in the restricted linear model when variance is non-constant.
- Fixed the alpha parameter initialization feature in the model options screen.
- New data grid windows now default to display 1000 rows. The previous default was 100.
- In addition to minor changes to accommodate the fixes and enhancements described above, the following Help file topics have been significantly revised or added:
  - Added: Troubleshooting>No Thousands Separator Can Be Used in the Data
  - Added: Troubleshooting>Requesting Support using eTicket
  - Added: Using BMDS 2.6.0>Keeping BMDS Up to Date
  - Added: Model Descriptions>Multi-tumor (MS_Combo) Model Description
  - Added a note to the Output From Models>Text Output from Models>Dichotomous Model Text Output topic, describing how standard error methods are calculated for parameters for multistage models vs. other dichotomous models. Also added text to the Continuous Model Text Output and Nested Model Text Output topics describing how their parameter standard errors are obtained.
  - Revised text on standard error calculations for the continuous model topics.
  - Updated the Model Descriptions>Nested Model Descriptions topic to describe how the goodness-of-fit p-values are calculated using a bootstrap approach. The Nested Model Description subtopics also include revised formula descriptions.
  - Updated the topic "Models Included in BMDS 2.6.0" to reflect updated versioning information.
  - Updated several topics in "Appendix A: Model Input File Format Descriptions" to reflect user interface changes. Also added input file details for the Dichotomous Hill and Quantal with Background Dose models.

- Updated several topics in "Appendix B: Model Options Screen Fields Reference" to reflect user interface changes.

A.20   BMDS 2.6.0.1

BMDS 2.6.0.1 contains the following fixes and enhancements:

- Resolves an issue in BMDS 2.6 that prevented users from specifying parameter values for Continuous models. (BMDS still prevents users from specifying the Continuous model Alpha parameter to work around an open issue where BMDS sometimes reports a non-optimal A3 log likelihood when the Alpha parameter is specified.)
- Resolves an issue with previous versions where BMDS occasionally reported incorrect scaled residuals when a session referenced multiple data files with differing dose values.
- Enhances the Nested models' text results to display the minimum, maximum, and mean of the absolute value of the scaled residuals, as well as the minimum, maximum, and mean of the scaled residuals.
- Resolves an issue where BMDS 2.6 could not display plots generated by previous BMDS versions.
- Implements a fix for an intermittent issue where the BMDS Wizard did not import plots generated by BMDS 2.4 and later.
- Correctly validates the v parameter for the Continuous Hill model, an issue since BMDS 2.3. Although the v parameter must be $0 < v \leq 1$ for Dichotomous Hill, it should be unconstrained for the Continuous Hill model.
- Adds an Update button to the About BMDS box (accessible from the Help>About menu item). Click the Update button to have BMDS check for and optionally install a BMDS update.

A.21   BMDS 2.7

August 18, 2017 - EPA distributed BMDS 2.7, which features several enhancements and fixes, including:

- The benchmark dose upper confidence limit (BMDU) is now included with the BMDL for most BMDS models (primarily continuous and dichotomous).
- BMDS installation now uses a Windows Installer msi-based file, rather than the zip files of previous releases. The .msi-based release is simpler, more robust, and minimizes problems related to Windows 10's increased security protocols.
- The auto-update feature (introduced in BMDS 2.6) works in a wider range of connection scenarios so users can be assured they are running the most recent version.

The Excel-based BMDS Wizard v1.11 is also included as part of the install package. BMDS Wizard remains unchanged except for minor fixes needed for compatibility with BMDS v2.7.

See the BMDS 2.7 readme file for details on this version's enhancements and fixes.

Appendix B: Citation Format and Acknowledgements

Citation Format

U.S. EPA (Environmental Protection Agency), 2018. Benchmark Dose Software (BMDS) Version 3.0. National Center for Environmental Assessment. Available from: https://www.epa.gov/bmds/benchmark-dose-software-bmds-version-30-download (Accessed mm/dd/yyyy).

Sponsors

EPA's National Center for Environmental Assessment with technical support from EPA's National Health & Environmental Effects Research Laboratory.

Contractor

General Dynamics Information Technology (GDIT)—BMDS user interface with a risk assessment modeling framework

Credits

- Dr. Matthew W. Wheeler, of The National Institute for Occupational Safety and Health (NIOSH), programmed both the continuous and dichotomous models for BMDS 3.0. His involvement consisted of rewriting all the models. This new code base allows for model averaging and future extensions to the BMDS modeling platform. He also developed the theory that allows the fast Bayesian computation of the model averaging results.
- RIVM (National Institute for Public Health and the Environment of the Netherlands, part of the Dutch administration), is a recognized leading center of expertise in the fields of health, nutrition, medicines, consumer safety and environmental protection working mainly for the Dutch government. RIVM is not a part of the US Environmental Protection Agency, but has given US EPA a non-exclusive, limited and revocable permission to use its logo only in recognition of RIVM's contribution to the development of certain BMDS models. The use of this logo here is neither an endorsement nor an advertisement for RIVM. RIVM does not accept any responsibility or liability for the activities of—or failures to act by—US Environmental Protection Agency.

We gratefully acknowledge the following software used in BMDS:
- Microsoft Excel 2016
- R

Note

Obtain the latest version of this program from the EPA BMDS web site. The features and models in the BMDS program that are downloadable from this web site have received at least an internal EPA review.

Disclaimer

This software has been reviewed in accordance with U.S. Environmental Protection Agency policy and approved for use. Mention of trade names or commercial products does not constitute endorsement.